NUSC Technical Report 5769

Measurements of Noise, Signal, And Signal-Plus-Noise Spectra Through a 5-Bit-Plus-Sign A/D Converter

John G. DePrimo
Design Department

D D C
MAR 10 1978
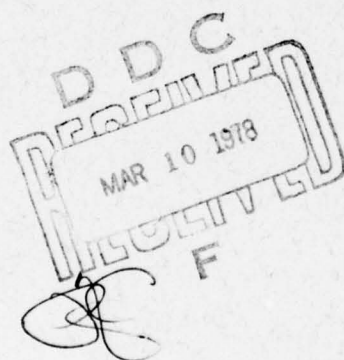
5 January 1978

NUSC

NAVAL UNDERWATER SYSTEMS CENTER
Newport, Rhode Island ● New London, Connecticut
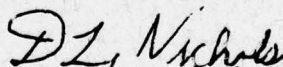
PREFACE

This report was prepared under the following programs:

1. The Ocean Measurements and Technology (OMAT) Program portion of the SEAGUARD program sponsored by the Defense Advanced Research Projects Agency (ARPA Order No. 2976), Program Manager, R. Cook, Tactical Technology Office; NUSC Project No. A69690, Principal Investigator, G. L. Assard (Code 313), Program Manager, B. Cole (Code 3104).

2. The TARP program sponsored by the Naval Sea Systems Command (SEA 06H2), Program Manager, T. Oliver; NUSC Project No. C68026, Principal Investigator, H. S. Newman (Code 314), Program Manager, J. A. Marsh (Code 323).

The Technical Reviewer for this report was G. L. Assard (Code 313).

The author wishes to acknowledge the technical assistance and direction offered by Mr. Assard.

REVIEWED AND APPROVED: 5 January 1978

D. L. Nichols

D. L. Nichols
Associate Technical Director
for Engineering and Technical Support

The author of this report is located at the New London Laboratory, Naval Underwater Systems Center, New London, Connecticut 06320.

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS
BEFORE COMPLETING FORM

1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER

NUSC — TR-5769

4. TITLE (and Subtitle)

MEASUREMENTS OF NOISE, SIGNAL, AND SIGNAL-PLUS-NOISE SPECTRA THROUGH A 5-BIT-PLUS-SIGN A/D CONVERTER.

5. TYPE OF REPORT & PERIOD COVERED

Technical rept. 2

6. PERFORMING ORG. REPORT NUMBER

7. AUTHOR(s)

John G. DePrimo

8. CONTRACT OR GRANT NUMBER(s)

9. PERFORMING ORGANIZATION NAME AND ADDRESS

Naval Underwater Systems Center
New London Laboratory
New London, CT 06320

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS

ARPA Order-2976
A69690   C68026

11. CONTROLLING OFFICE NAME AND ADDRESS

Defense Advanced Research Projects Agency
Arlington, VA 22209
Naval Sea Systems Command (SEA 06H2), Wash, DC 20362

12. REPORT DATE

5 January 1978

13. NUMBER OF PAGES

51

14. MONITORING AGENCY NAME & ADDRESS(If different from Controlling Office)

53 p.

15. SECURITY CLASS. (of this report)

UNCLASSIFIED

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

DDC
MAR 10 1978
RECEIVED
F

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

A/D Converters
Quantizers
Quantized Spectra

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report presents the spectra of various quantized input signals obtained by laboratory measurement at the output of a 5-bit-plus-sign analog-to-digital (A/D) converter. The results demonstrate the spectral dynamic range capability for single and multitonal inputs, with and without injected band-limited noise. Optimum input level adjustment, that is, A/D converter biasing, is determined as a function of input signal type. These measurements are verified in theory, wherever possible.

DD FORM 1473
1 JAN 73

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF ILLUSTRATIONS (CONT'D)

GLOSSARY

Symbol	Interpretation
A	Normalizing constant less than or equal to 1
b	Number of bits
DC	Direct Current (zero hertz)
E	Amplitude of a sinusoid
K	An integer constant
N_q	Power in quantization process alone
P_n	Noise power resulting from quantizing a normalized random signal
p(x)	Probability density function of the random signal
r	The ratio of maximum input signal, total voltage swing to the voltage swing of one quantization step
t	Time
ω	Radian frequency
x(t)	A random signal
$\hat{x}(t)$	A quantized, random signal
ε_x	The absolute error generated due to the quantization process
σ_ε	Quantization process standard deviation (based on quantizer step size)
σ_ε^2	Quantization process variance (based on quantizer step size)
σ_{ε_x}	Standard deviation of the error due to quantization
$\sigma_{\varepsilon_x}^2$	Variance of the error due to quantization
σ_x	Standard deviation of a random signal
σ_x^2	Variance of a random signal
μ_{ε_x}	The mean error due to quantization, assumed to be zero in this report

GLOSSARY (CONT'D)

Symbol	Interpretation
A/D	Analog-to-digital converter
D/A	Digital-to-analog converter
dB	Decibels
dBV	Decibels referenced to 1 volt
Hz	Hertz
kHz	Thousands of hertz
LPF	Lowpass filter
LSB	Least significant bit
M_{BL}	Noise bandlevel
M_{equiv}	Equivalent noise bandlevel
M_{quant}	Quantization noise level
mV	Millivolt (10^{-3})
N	Noise
rms	Root mean square
S	Signal
S/H	Sample and hold
SNR	Signal-to-noise ratio
V	Volts
Vrms	rms volts

MEASUREMENTS OF NOISE, SIGNAL, AND SIGNAL-PLUS-NOISE SPECTRA THROUGH A 5-BIT-PLUS-SIGN A/D CONVERTER

INTRODUCTION

Laboratory measurements of analog-to-digital (A/D) converter output spectra for various inputs were conducted to confirm the practical performance limits attainable by using only 5 bits plus sign. Absolute quantization noise levels and signal-to-noise ratios (SNRs) can be derived analytically, and in practice they are approximated by very simple rules of thumb. In general, hardware measurements verify the theoretical parameters, and may also point out pitfalls due to the nonideal nature of actual A/D converters. The measurements would also quantify the sensitivity of these analytical predictions to proper amplitude biasing at the quantizer input. This particular information is essential for the proper definition of automatic or time-varying gain control characteristics, and their effects on realizable system performance.

HARDWARE

Figure 1 illustrates the basic test configuration used for the measurements. The equipments illustrated in the figure may or may not be used in each test but are presented for background information. A Micronetworks Corporation, 12 bit A/D converter and analog sample and hold (S/H) device are used at the input. The lower six bits are masked off, which results in doubling the effective step size for every bit masked, and reduces the zero to positive sign-bit threshold by 1/64 relative to the redefined least significant bit (LSB). The sign bit appears to hard clip zero-to-positive or negative-to-positive transitions much smaller than the redefined LSB step size, preserving phase information when amplitude information is lost. The digital data are latched for subsequent digital-to-analog (D/A) conversion by a Hybrid Systems DAC 331. This is an eight bit device, with the two lower-order bits masked. Additional errors introduced by either the DAC or the S/H device, such as slewing, glitching, or data skewing, contribute little to the results presented here. These functions are routinely associated with an A/D converter; hence, all errors are considered inseparable when conducting laboratory measurements. Both the A/D and the D/A use offset binary code.

A 3 kHz sample rate was required at the S/H input. With the test circuit configured as shown, the master clock frequency had to be 13 times 3 kHz (39 kHz) to allow for the completion of the successive approximation of the A/D converter. On several occasions, it was necessary to increase the sample rate by a factor of ten to demonstrate a specific effect.

Some of the lowpass filtered data presented in this report have low frequency rolloff due to capacitive coupling at the quantizer input.
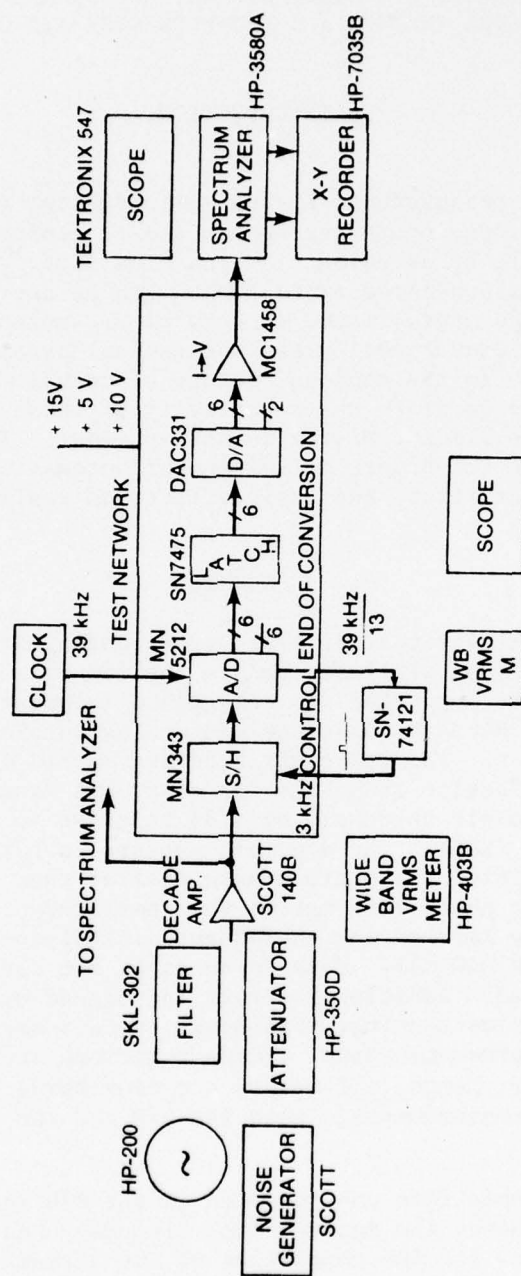
Figure 1. Basic Test Configuration

SINGLE-TONE INPUT

The first set of measurements was performed with a single, multilevel tone at the input. Maximum SNR for a single tone is realized if the tone fully occupies all the converter bits while maintaining an acceptable level of harmonic distortion. This condition can be assured by monitoring the unfiltered D/A output with both an oscilloscope and a spectrum analyzer. Since this particular A/D converter has a specified range of 20 V peak to peak, 7 Vrms was chosen as the full tone (maximum) input signal. The harmonic content evident in figures 2 and 3 proved that the signal was distorted, although this was not obvious when examining the time waveform on the oscilloscope. The input was stepped down 2 dB at a time, until there was no noticeable decrease in harmonic energy for a corresponding decrease in input level. A 4 dB reduction at the input yielded a 10 dB increase in harmonic suppression. However, figures 2 through 4 show little variation in the actual noise floor. In applications where the system passband is narrow enough to exclude the harmonic frequencies, the converter can apparently be driven slightly into saturation and still realize a small increase in SNR. Figures 5 through 9 show the converter output spectrum for input levels reduced 10 dB at a time. Statistical modeling of an A/D converter is invariably done assuming a Gaussian input, which assures that the probability distribution of ε_x is uniform, as noted. We expect, therefore, that the available theoretical predictions of SNR applied to a single tonal are not precise. It is shown later in this section that the harmonics generated by the quantization process do indeed confound both predictions and measurement. Theoretically, SNR can be calculated as

$$SNR = 10 \log_{10} \frac{A^2 \sigma_x^2}{\sigma_\varepsilon^2} \; ,$$

where σ_x^2 and σ_ε^2 are signal and noise (quantization error) variances, respectively. The input signal $x(t)$ is normalized by A, such that $Ax(t)$ does not exceed the full-tone SNR capability of the A/D converter. This SNR expression is valid over the interval from dc to the Nyquist frequency (one-half the sampling frequency), where quantization noise is considered uniformly distributed. In practice, actual SNR is a function of sample rate and of system processing gain due to operating bandwidths less than the Nyquist frequency. The higher the ratio of Nyquist frequency to operating band, the greater the achievable SNR.

Since the input signal, in this case, is a sinusoid, it is considered statistically stationary. More specifically, it is an ergodic process. It can therefore be written as

$$\sigma_x = x(t) \text{ Vrms,}$$

where

$$x(t) = E \sin \omega t.$$

Therefore,

$$\sigma_x = E/\sqrt{2} \; .$$

The quantization error ε_x is defined as the difference between the exact signal value and its quantized value

$$\varepsilon_x = |\hat{x}(t) - x(t)|,$$

where $\hat{x}(t)$ is the quantized version of $x(t)$.

This equation states that ε_x can attain values between $\pm 1/2$ LSB or a normalized peak-to-valley error of 2^{-b}, where b is the number of bits. The mean and variance of ε_x are calculated as follows:[1]

The mean, assuming that rounding between converter steps occurs and that the probability density function of ε_x is uniform, is

$$\mu_{\varepsilon_x} = \int_{-1/2 \text{ LSB}}^{+1/2 \text{ LSB}} \varepsilon_x d\varepsilon_x = 0.$$

The variance, with the above assumptions, is

$$\sigma_\varepsilon^2 = E[(\varepsilon_x - \mu_{\varepsilon_x})^2] = \int_{-1/2 \text{ LSB}}^{+1/2 \text{ LSB}} \varepsilon_x^2 d\varepsilon_x = \frac{2^{-2b}}{12}$$

and $\sigma_{\varepsilon_x} = 2^{-b}/\sqrt{12}$ is the standard deviation.

Maximum, full-tone SNR is defined by the same signal and noise variances, except that the ratio $2E/\varepsilon_x = r$ is established.

The total range of the quantizer must be 2E, since we have defined this as full tone. The constant r, then, is the ratio of one quantization step to the total voltage range of the A/D converter.[2]

Since

$$\frac{\sigma_x^2}{\sigma_\varepsilon^2} = \frac{E^2/2}{E^2/12} \qquad \text{and} \qquad \frac{2E}{\varepsilon} = r,$$

then

$$\frac{\sigma_x^2}{\sigma_\varepsilon^2} = \frac{3r^2}{2} \quad .$$

4

Therefore, $10 \log\left(3r^2/2\right)$ is the predicted full-tone SNR in a noise band from dc to the Nyquist frequency. For this 5-bit-plus-sign converter,

$$\text{Full-tone SNR} = 10 \log \frac{3(64)^2}{2} = 37.9 \text{ dB (rms)}.$$

In a 1 Hz band,

$$\text{Full-tone SNR} = 37.9 + 10 \log 1500 = 69.7 \text{ dB} \quad \text{(rms)}.$$

To fully appreciate the physical significance of the full-tone SNR prediction, the following example using a 5-bit-plus-sign converter is offered:

Let $\varepsilon = 1$ step (smallest step size).

Let $E = 5$ magnitude bits or 32 steps.

Then

$$\frac{E^2/2}{\varepsilon^2/12} = \frac{(32)^2/2}{1/12} = (32)^2 \times 6 = 6144, \quad 10 \log 6144 = 37.88 \text{ dB, as}$$

above.

This number is verifiable in figure 2, within a decibel or so. Note that, on a log scale, the noise mean amplitude is not found by drawing a line through the middle of the plot. The Hewlett-Packard spectrum analyzer used here has selectable smoothing times which reduce the noise variance, but result in excessively long plotting times.

To solve for the normalizing constant A in the specific case of a full input tone, calculate σ_ε and $\sigma_x(\text{max})$ from the available information:

$$\sigma_\varepsilon = \frac{1/64}{\sqrt{12}} = 4.511 \times 10^{-3} \quad \text{and} \quad \frac{E}{\sqrt{2}} = 7 \text{ Vrms}$$

$$20 \log A = \text{full-tone SNR (dc to Nyquist)} - 20 \log 7 + 20 \log 4.511 \times 10^{-3}$$

$$20 \log A = -25.9 \text{ dBV}$$

$$A = 0.05.$$

Therefore, the analytical expression predicting the SNR for this converter, from dc to the Nyquist frequency, is

$$\text{SNR} = 10 \log_{10} \frac{0.05\sigma_x^2}{\sigma_\varepsilon^2} \; ;$$

so

$$SNR = (20 \log \frac{E}{\sqrt{2}} + 20.9) \text{ dB}$$

and

$$SNR = (20 \log \frac{E}{\sqrt{2}} + 52.7) \text{ dB}/\sqrt{Hz}.$$

The noise floor of the A/D converter, normalizing the peak level to 1, is defined simply by the quantization noise power,

$$N_q = 10 \log \sigma_\epsilon^2 \quad , \quad \text{dc to Nyquist}$$

$$N_q = -40.9 \text{ dBV peak}$$

$$N_q = -72.7 \text{ dBV peak}/\sqrt{Hz}.$$

This is the level predicted based upon the step size alone, independent of any input. It is not a practical design limit. Note that the difference between the theoretical, normalized spectral noise floor and the full tone $(SNR)^{-1}$ will always be 9 dB, independent of the number of bits.

Figure 10 is a graph of predicted and measured SNR for this A/D converter. The input level reaches a lower limit where the LSB or the sign bit fails to produce meaningful zero-crossing information. This level appears to be slightly less than -40 dBV rms.

Examination of the tonal levels in figures 2 through 9 will be helpful in recognizing the region of linear A/D converter operation indicated in figure 10. At +17, +15, and +13 dBV input levels, the A/D output spectral levels do not quite indicate a 2 dB difference between tonals. It is evident that the converter is still slightly in saturation at +13 dBV input, although the harmonic levels are acceptable. At 0, -10, and -20 dBV input levels, the output tonals follow the input, and decrease 10 dB each time. Linear operation is realized(experimentally verified)between these input levels. The input level change from -20 to -30 dBV produces only a small change at the output, indicating the nonlinear operating condition expected when only the sign bit and/or the LSB are being activated by the input. At -40 dB Vrms, the signal output as monitored on the oscilloscope is aperiodic, asymmetric, and positive going (see the section on HARDWARE). Occasional bursts of information containing the signal zero crossings come through, and the spectral line is recognizable (see figure 9).

The converter is slightly driven into saturation at levels around +13 dBV. This condition is recognized in figure 10 by the change in measured SNR slope. An insufficient number of data points were taken to accurately define that portion of the curve between 0 and +13 dBV. It is certain that linear operation extends somewhat above 0 dBV input.

It is clear that a 6 dB increase in SNR over the predicted value is realized within the linear operating region of the converter. This increase is possibly caused by the inaccuracy of an "eyeball" estimation of the noise mean, which is ambiguous in the case of a pure tonal due to the generation of many spectral lines above the apparent noise floor. Refer to the quantized, 200 Hz band-limited noise curves of figure 11A. The second curve corresponds to an absolute input (wideband rms average) of +1 dBV. This level is within a decibel of the tonal level presented in figure 5. To arrive at a comparison between the two SNRs, it is necessary to take the 200 Hz wideband noise level and convert it to an equivalent 1 Hz tonal (Analyzer bandwidth = 30 Hz):

$$M_{equiv} = M_{BL} - 10 \log 30 + 10 \log 200 \cong M_{BL} + 8.25 \text{ dB},$$

where M_{equiv} is the equivalent magnitude of the noise if it were concentrated into a 1 Hz band and M_{BL} is the noise passband level in the 200 Hz baseband illustrated in figure 11A as -13 dBV. This figure is discussed in detail later. Hence, $M_{equiv} = -4.75$ dBV.

The stopband noise floor in figure 11A is actually the quantization noise floor. Its level in a 1 Hz band is

$$M_{quant} = -43 \text{ dBV} - 10 \log 30 = -57.75 \text{ dBV}.$$

Therefore, the equivalent SNR is

$$M_{equiv} - M_{quant} = 53 \text{ dBV}.$$

This is within 1/2 dB of the SNR predicted by figure 10. It is therefore concluded that the SNR measured in the linear operating region of the converter is indeed optimistic, and that the noise mean should have been estimated about 6 dB higher.

Since the A/D converter analyzed here is an odd quantizer, that is, the converter transfer function is odd,[3,4] one would expect the device to produce only odd harmonics.[5] Figures 2-9 show that an occasional, even harmonic is clearly present. In addition, some higher-order harmonics have more energy than the third harmonic (see figure 6). Any combination of aliasing, harmonic, intermodulation, power supply harmonic, and sampling frequency tonals can produce results which deviate from those theoretically predicted. Maintaining a near-zero dc offset is also a problem, in that a dc spectral component is produced and distortion on one half-cycle (+ or -) is realized before the other. This produces even harmonics.
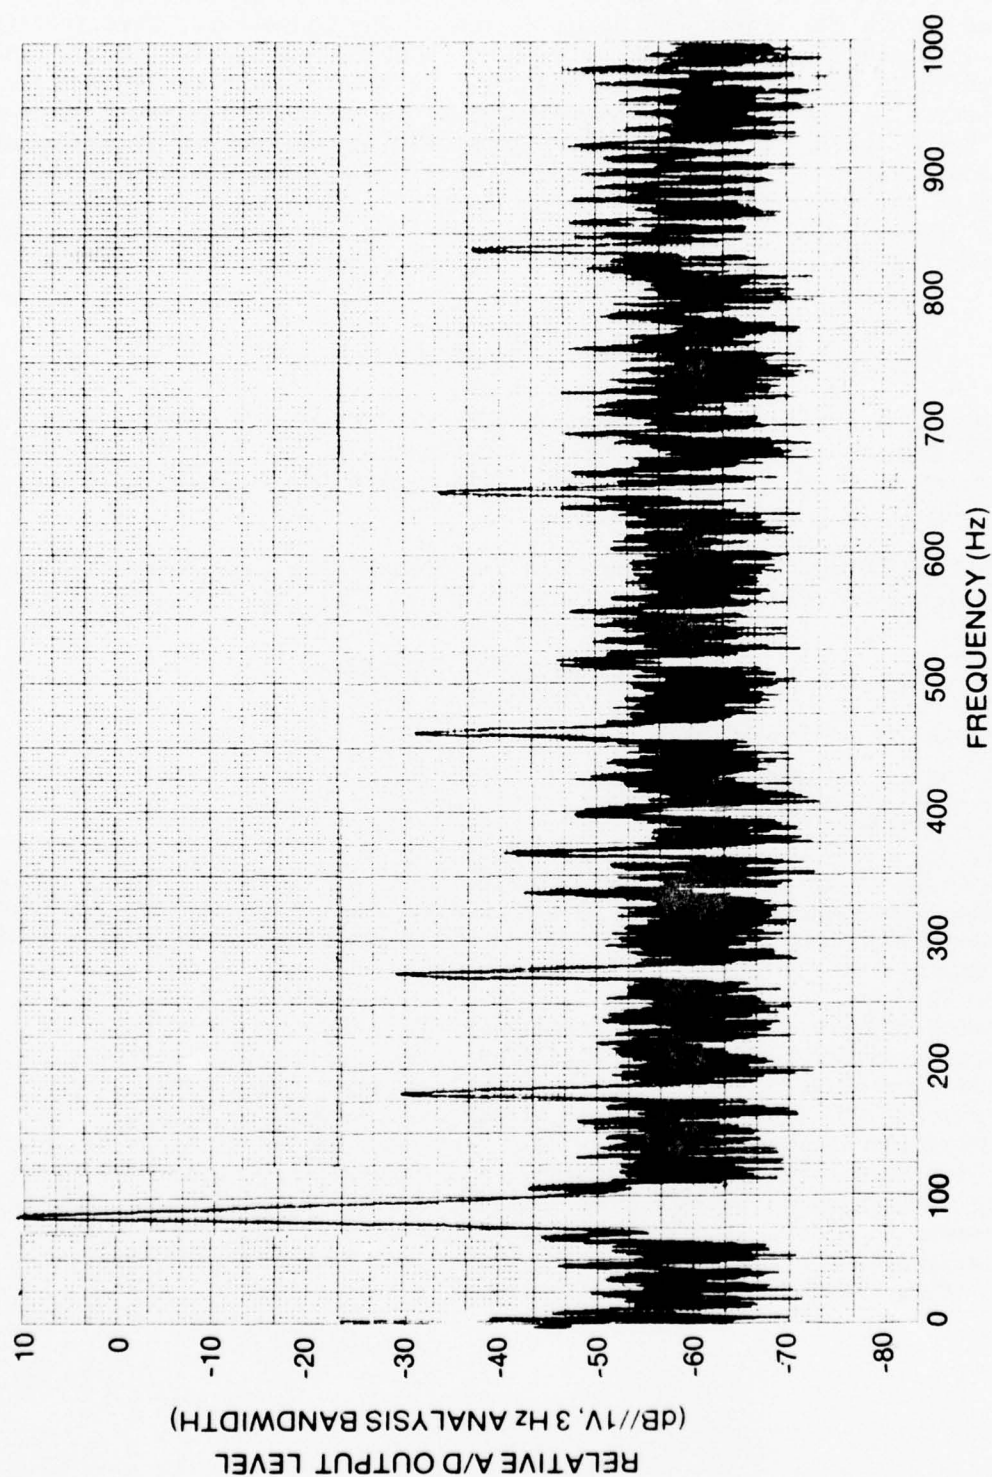
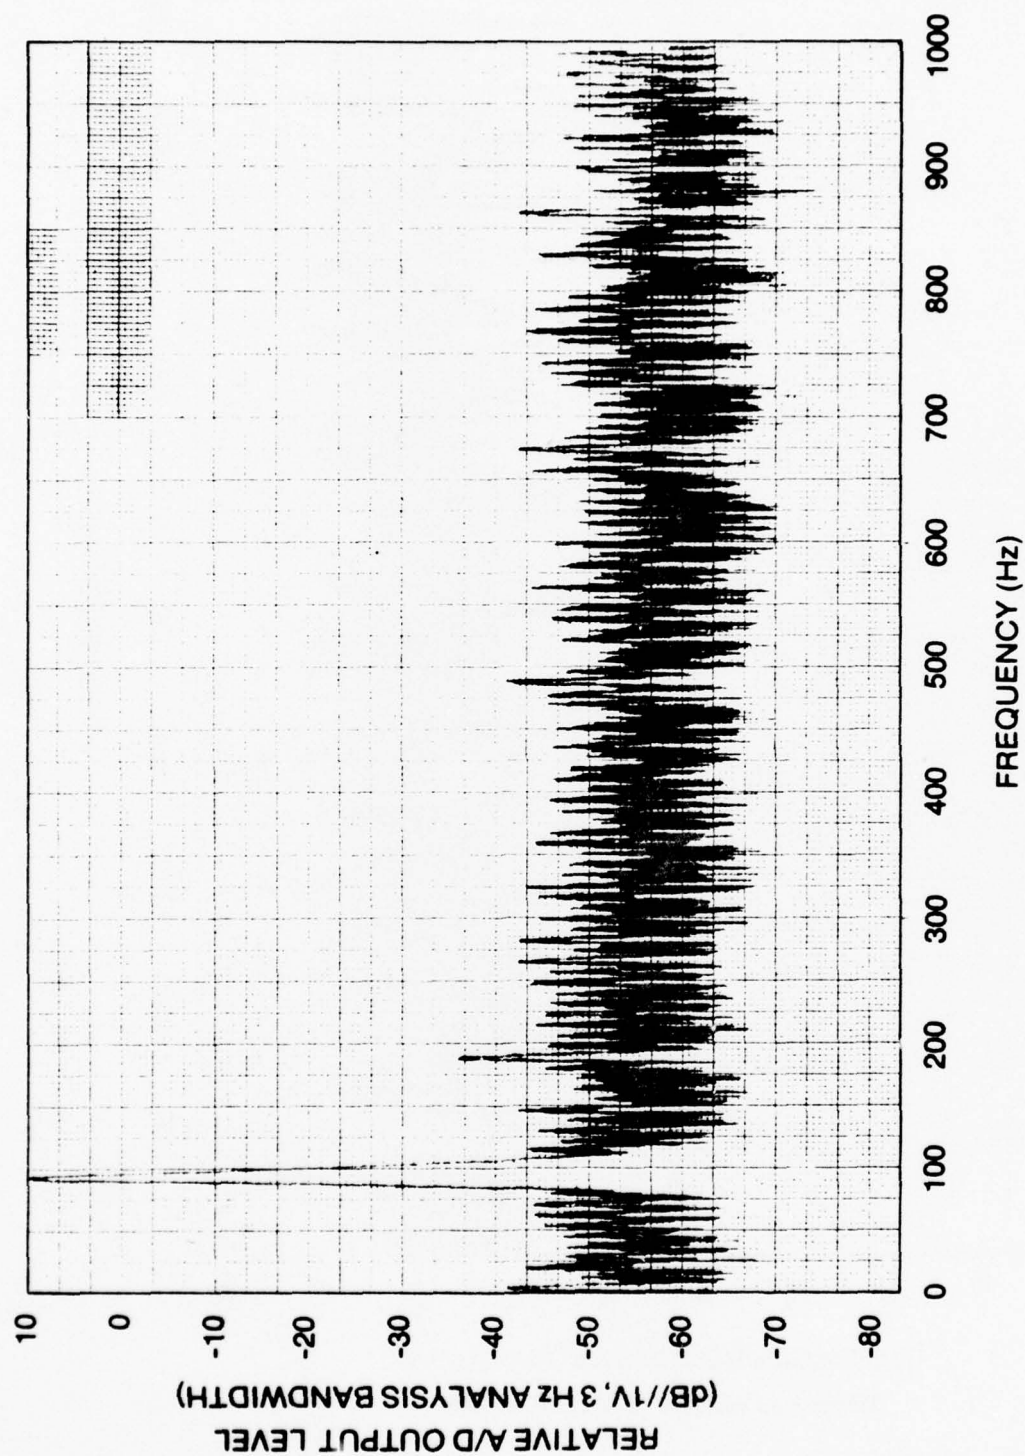Figure 2. Single Tone, +17 dBV Input Level

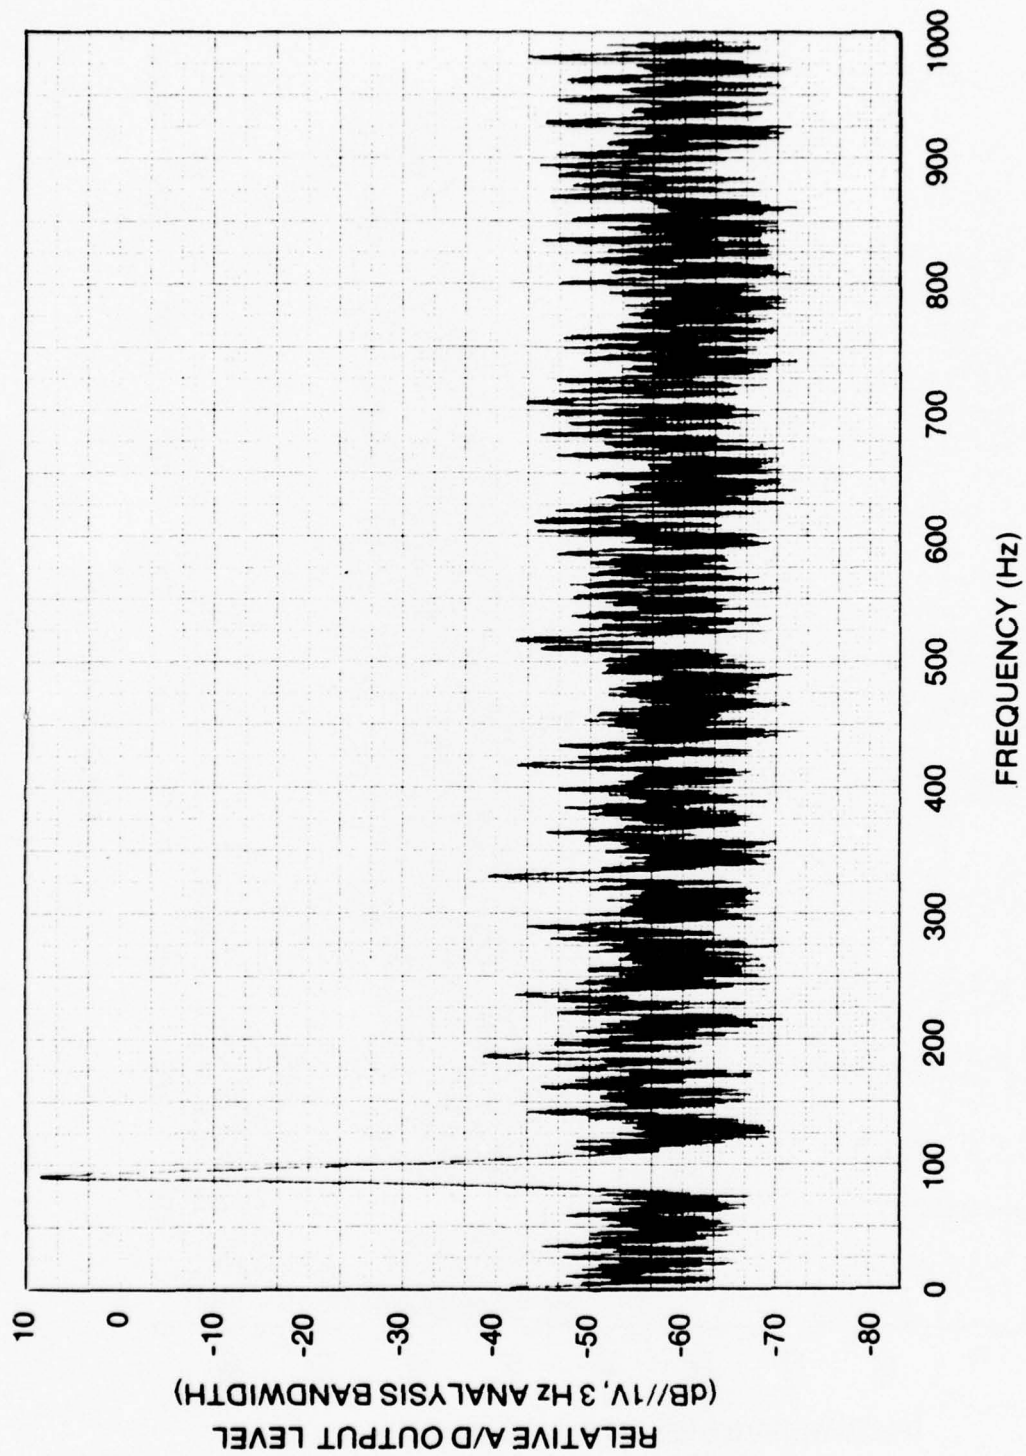Figure 3. Single Tone, +15 dBV Input Level
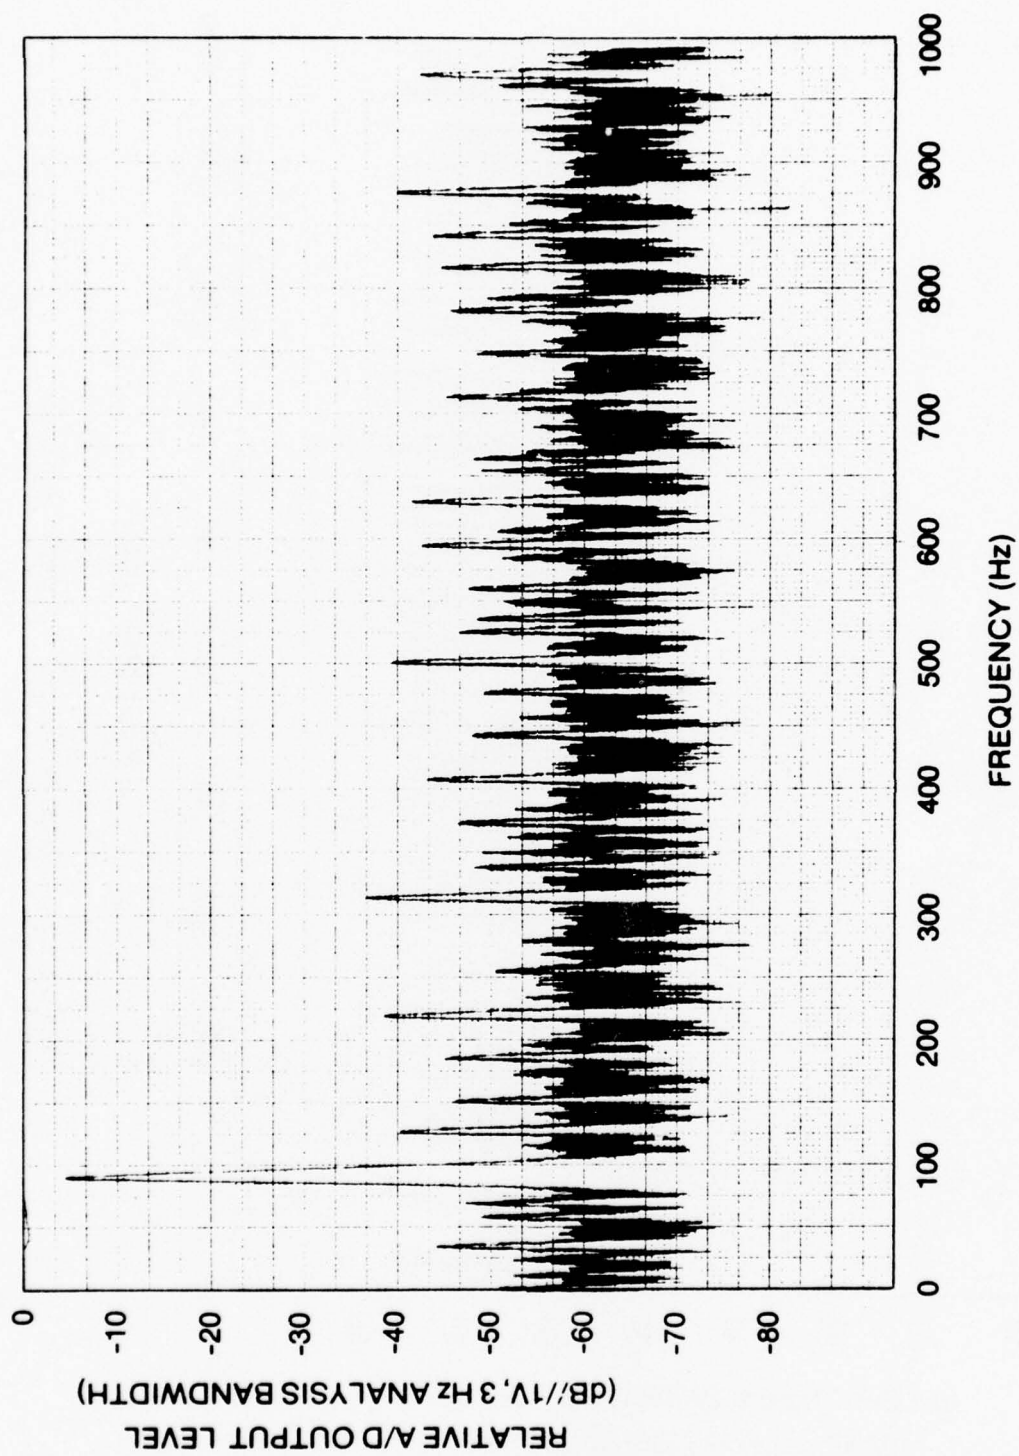
Figure 4. Single Tone, +13 dBV Input Level
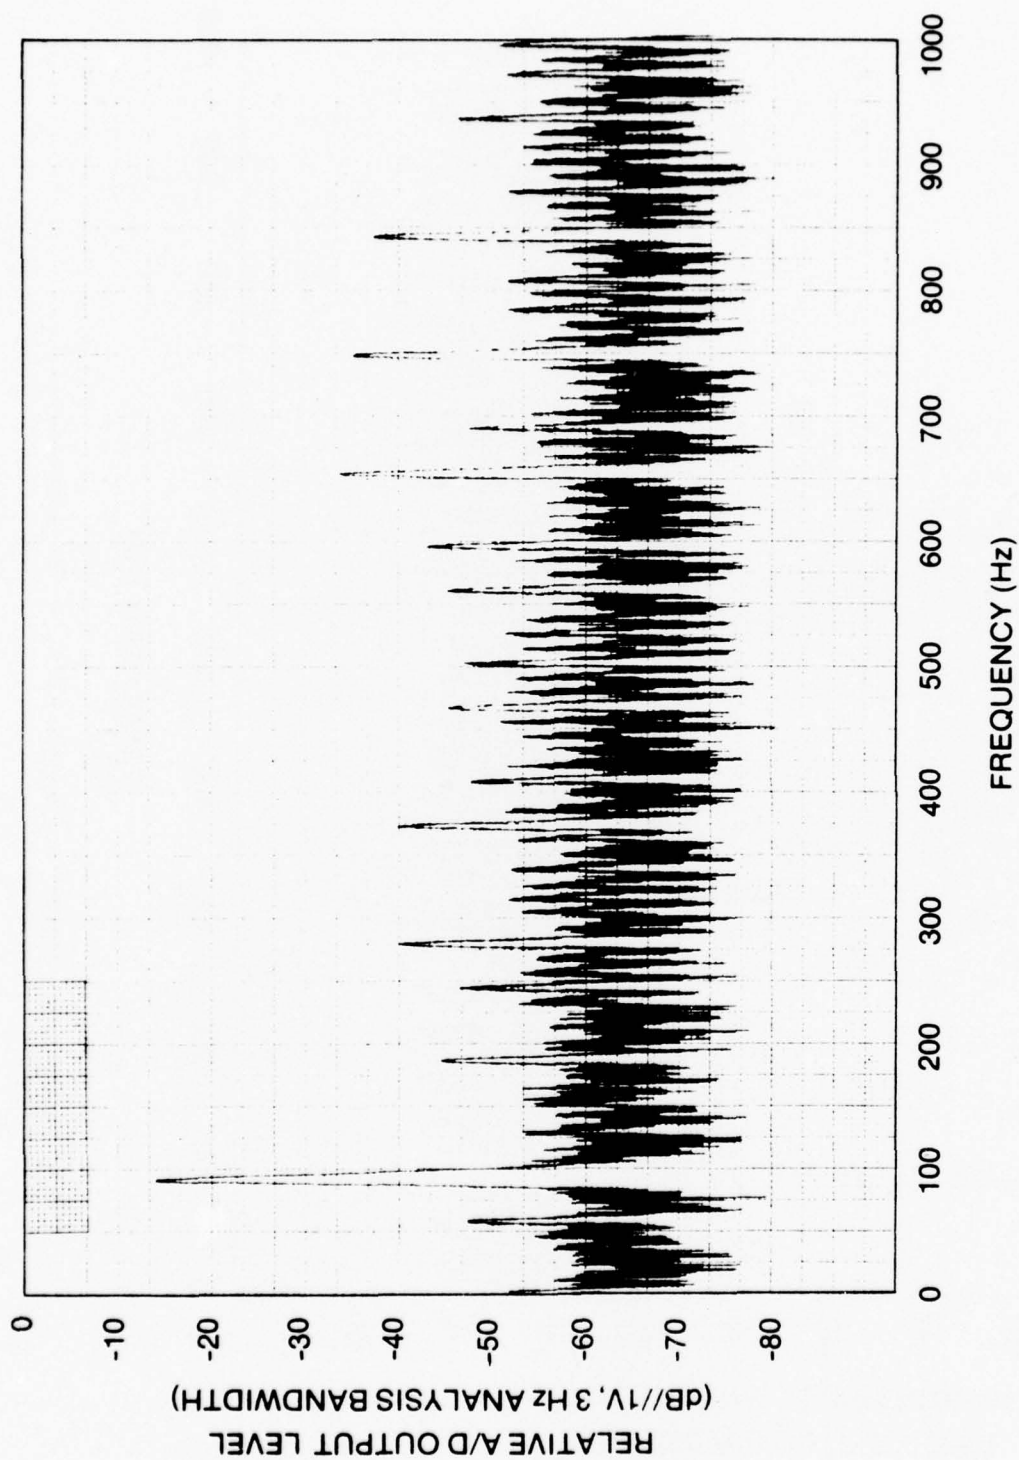
Figure 5. Single Tone, 0 dBV Input Level

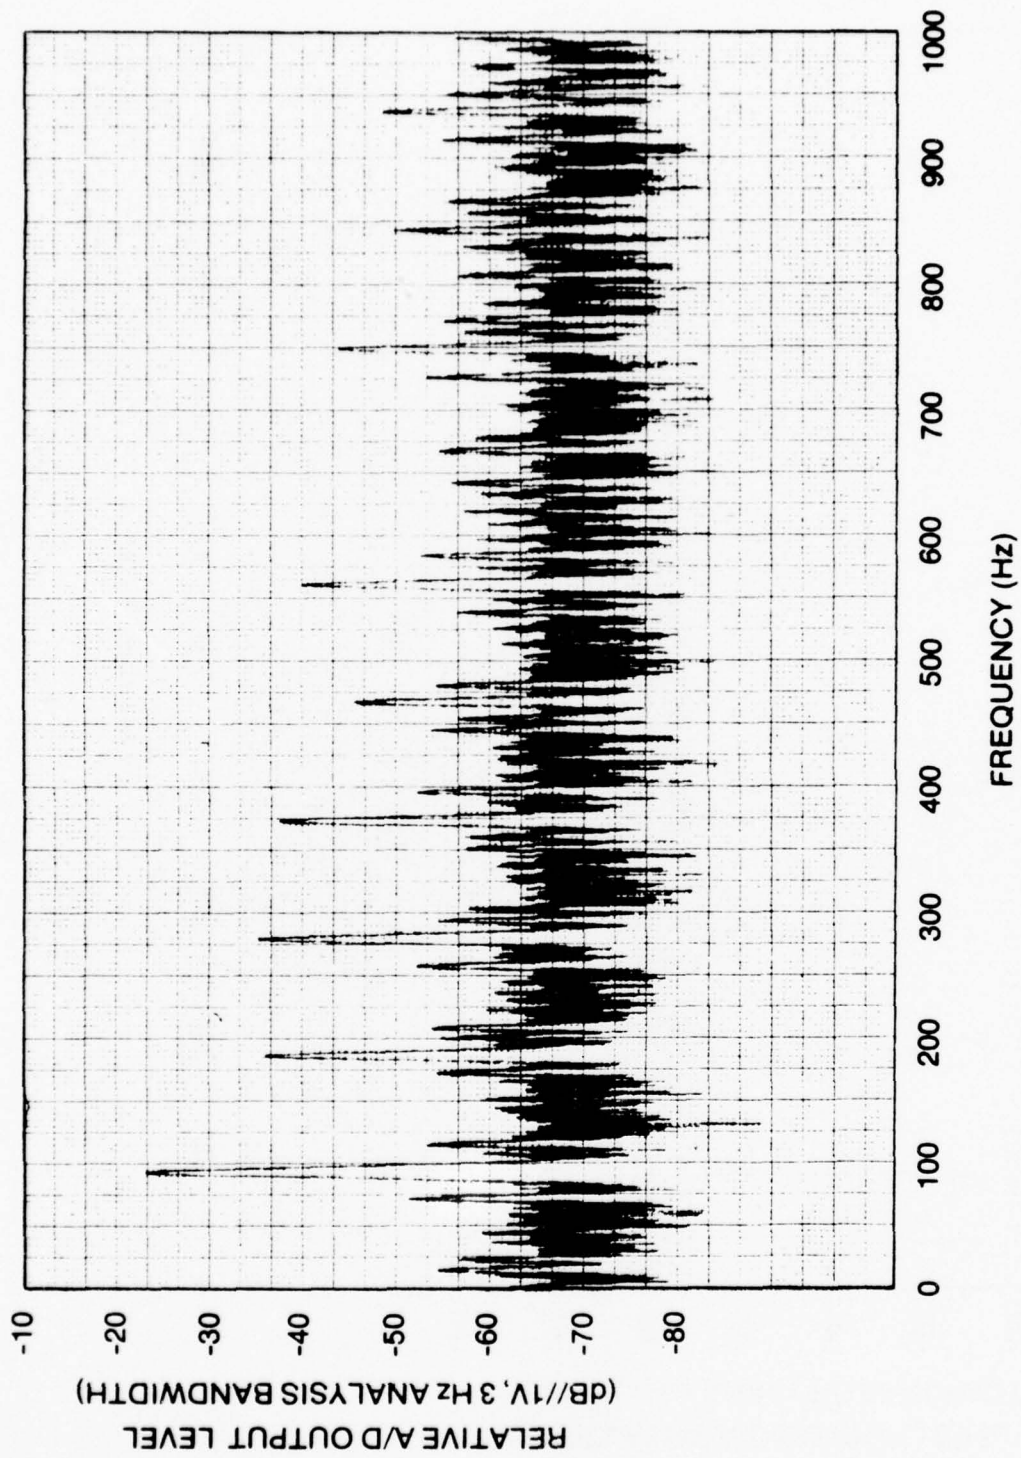Figure 6. Single Tone, -10 dBV Input Level

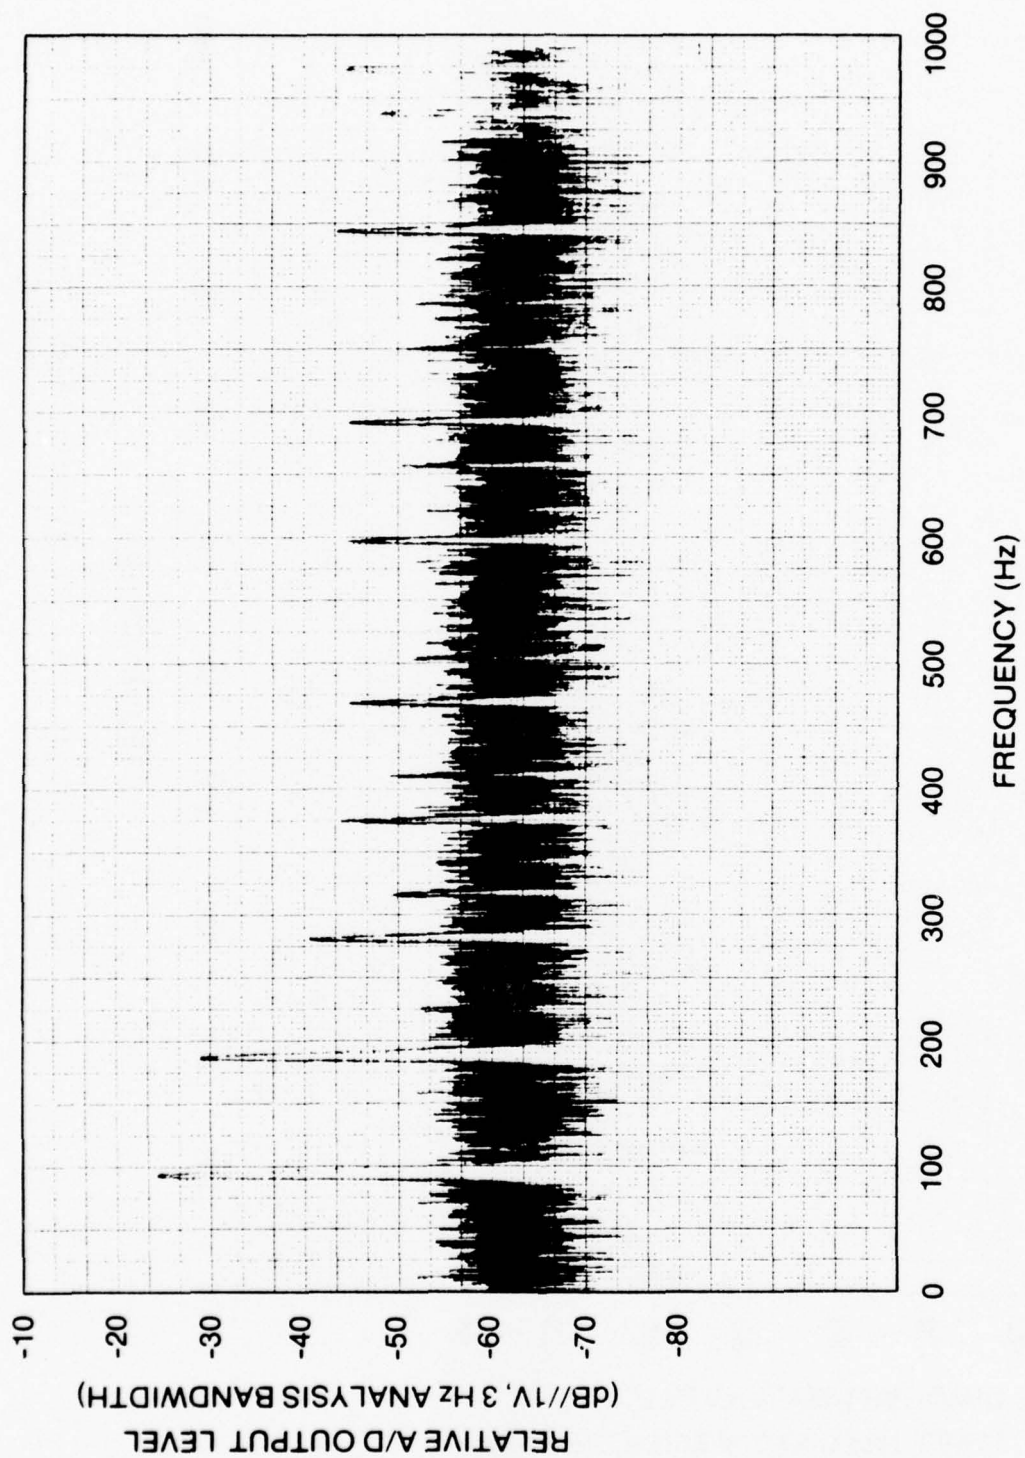Figure 7. Single Tone, -20 dBV Input Level

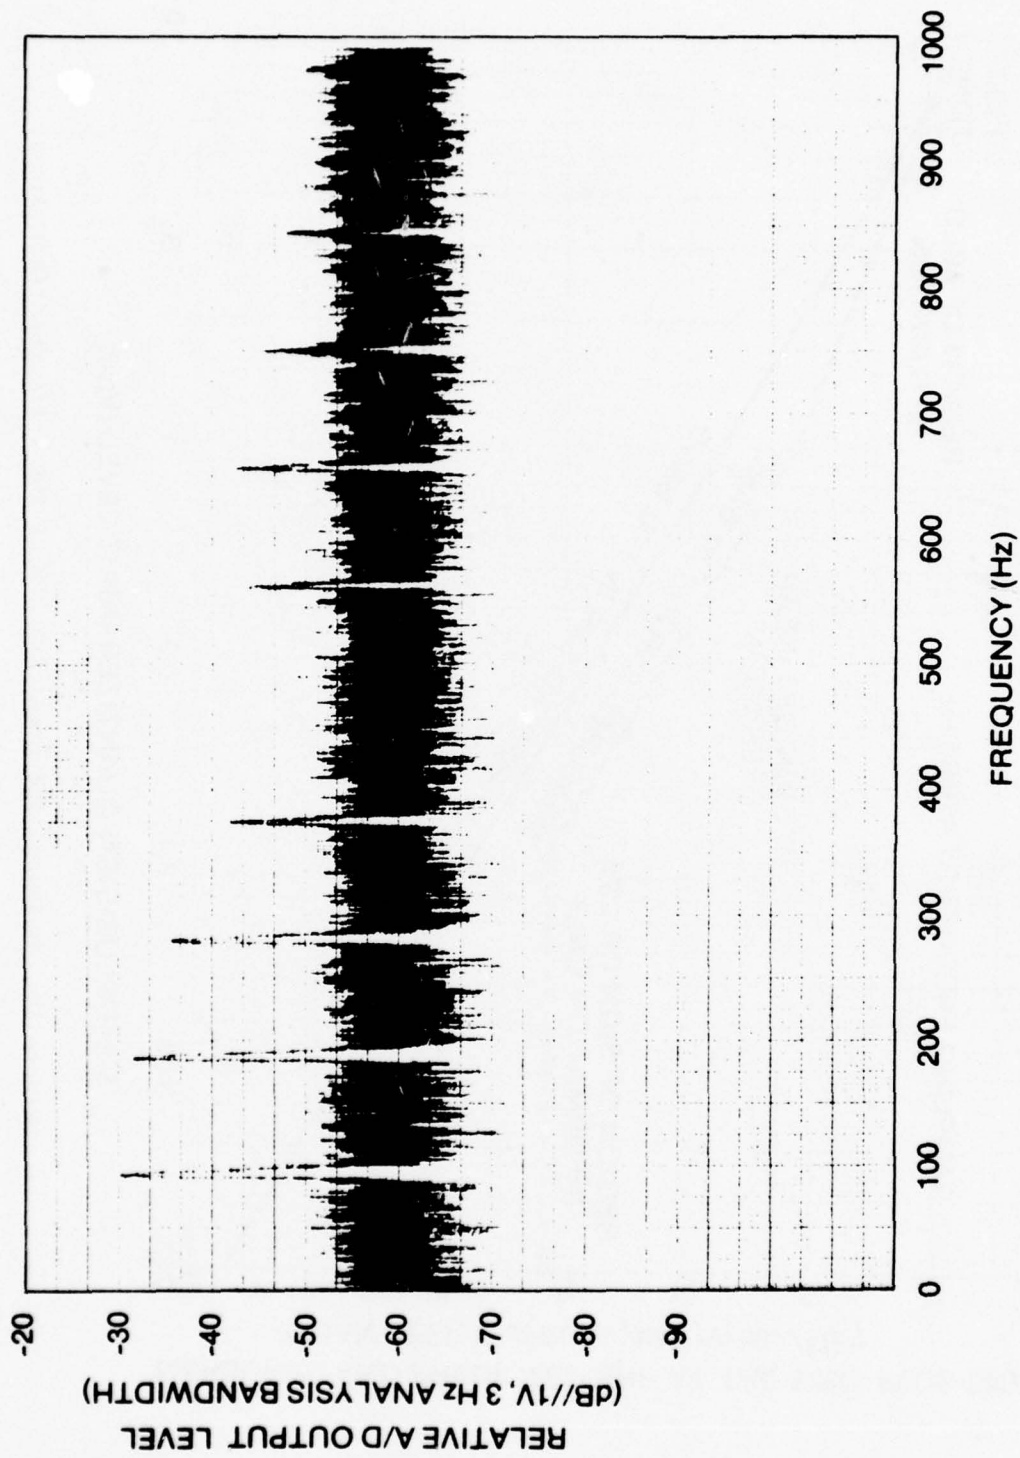Figure 8. Single Tone, -30 dBV Input Level

Figure 9. Single Tone, -40 dBV Input Level
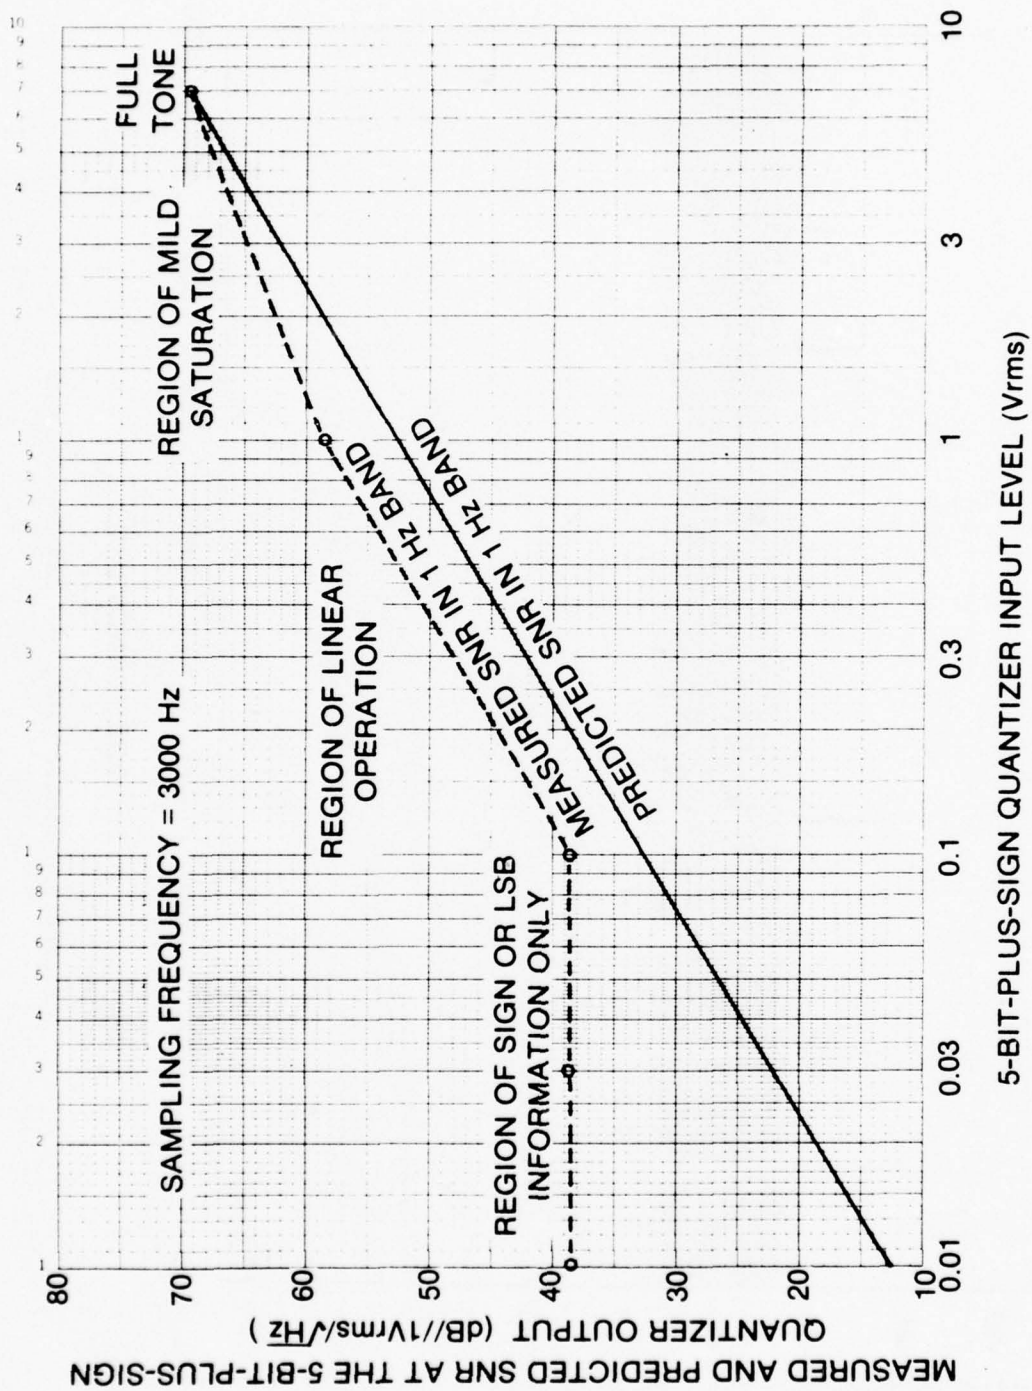
Figure 10. Measured and Predicted SNR at the 6-Bit Quantizer Output
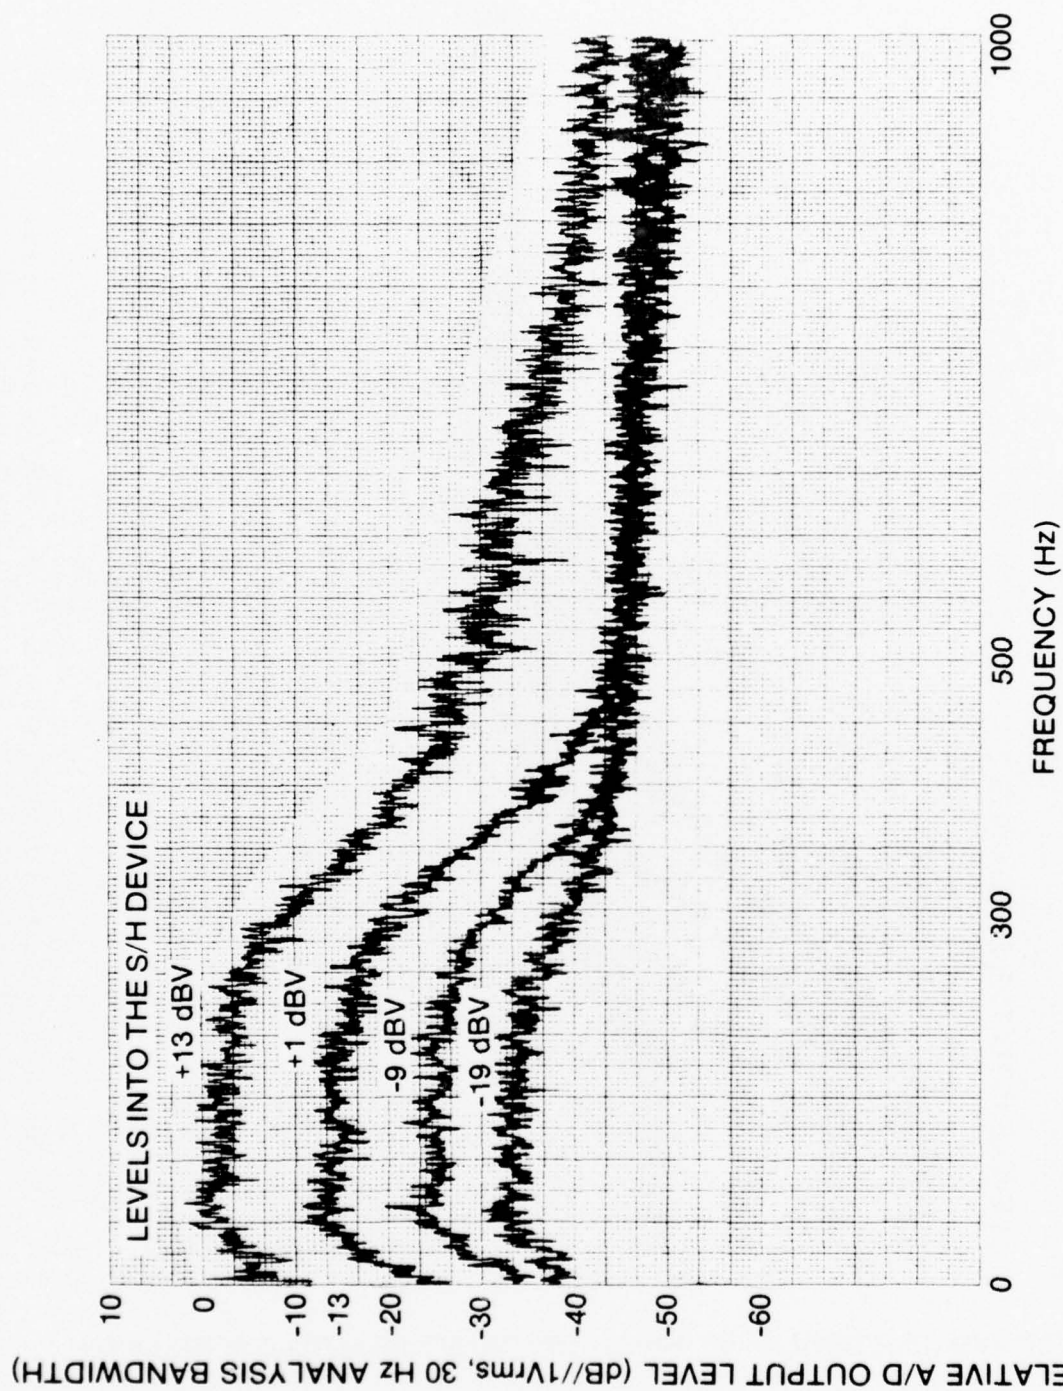
Figure 11A. Quantized Lowpass Filtered Noise, 300 Hz Corner Frequency, 3.0 kHz A/D Sample Rate, Input Levels as Indicated
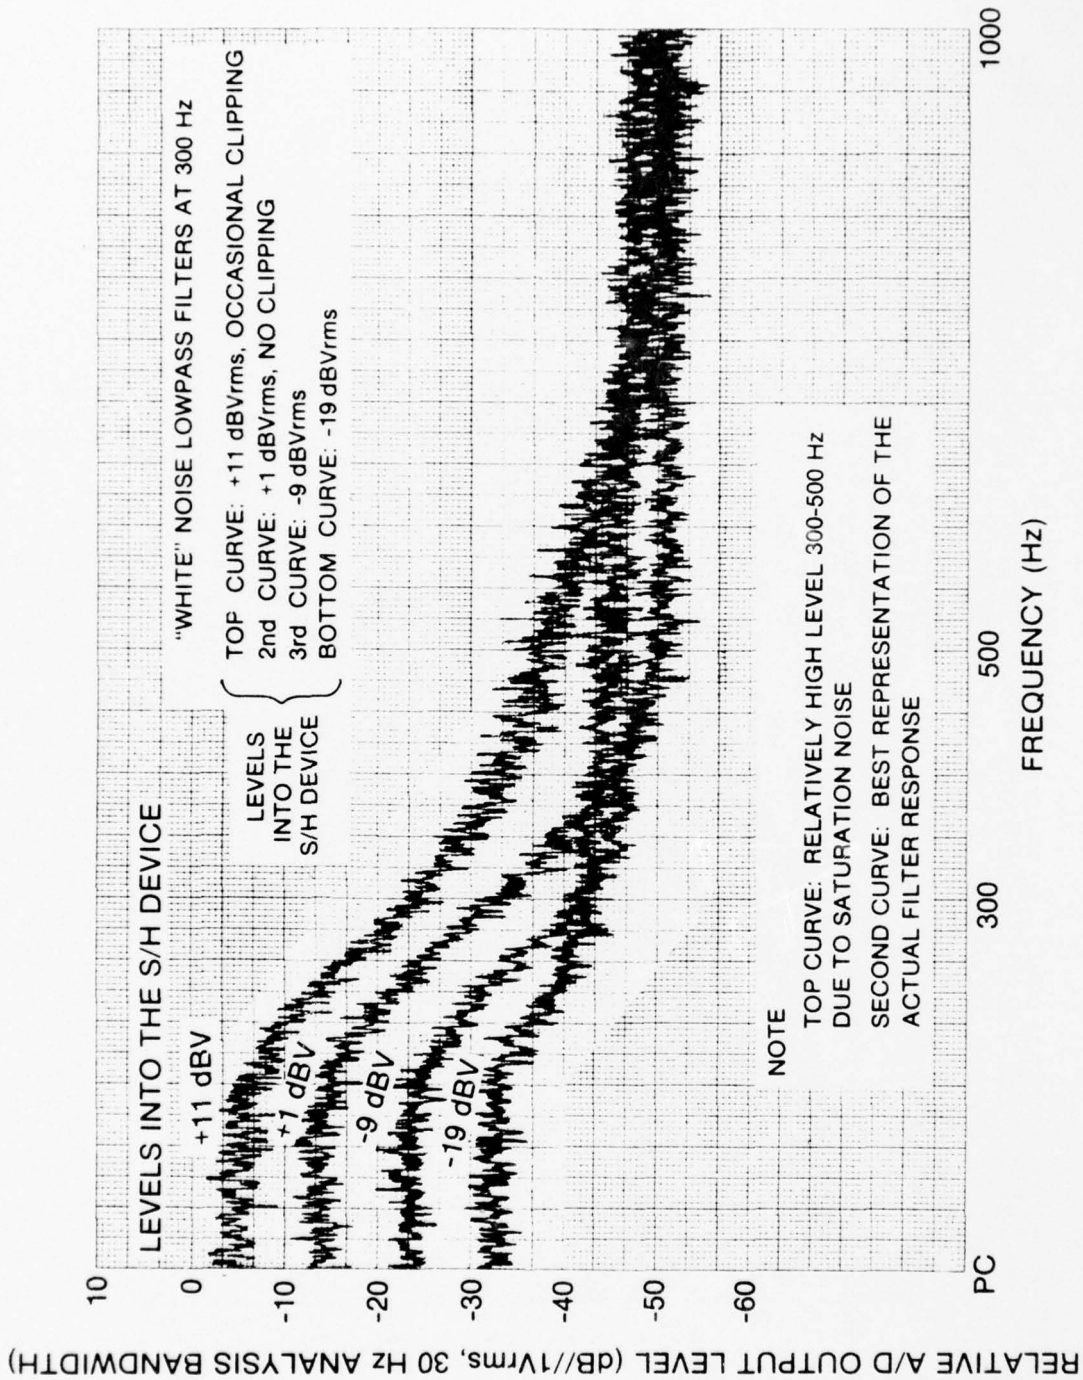
Figure 11B. Quantized Lowpass Filtered Noise, 300 Hz Corner Frequency, 30.0 kHz A/D Sample Rate, Four Input Levels as Indicated
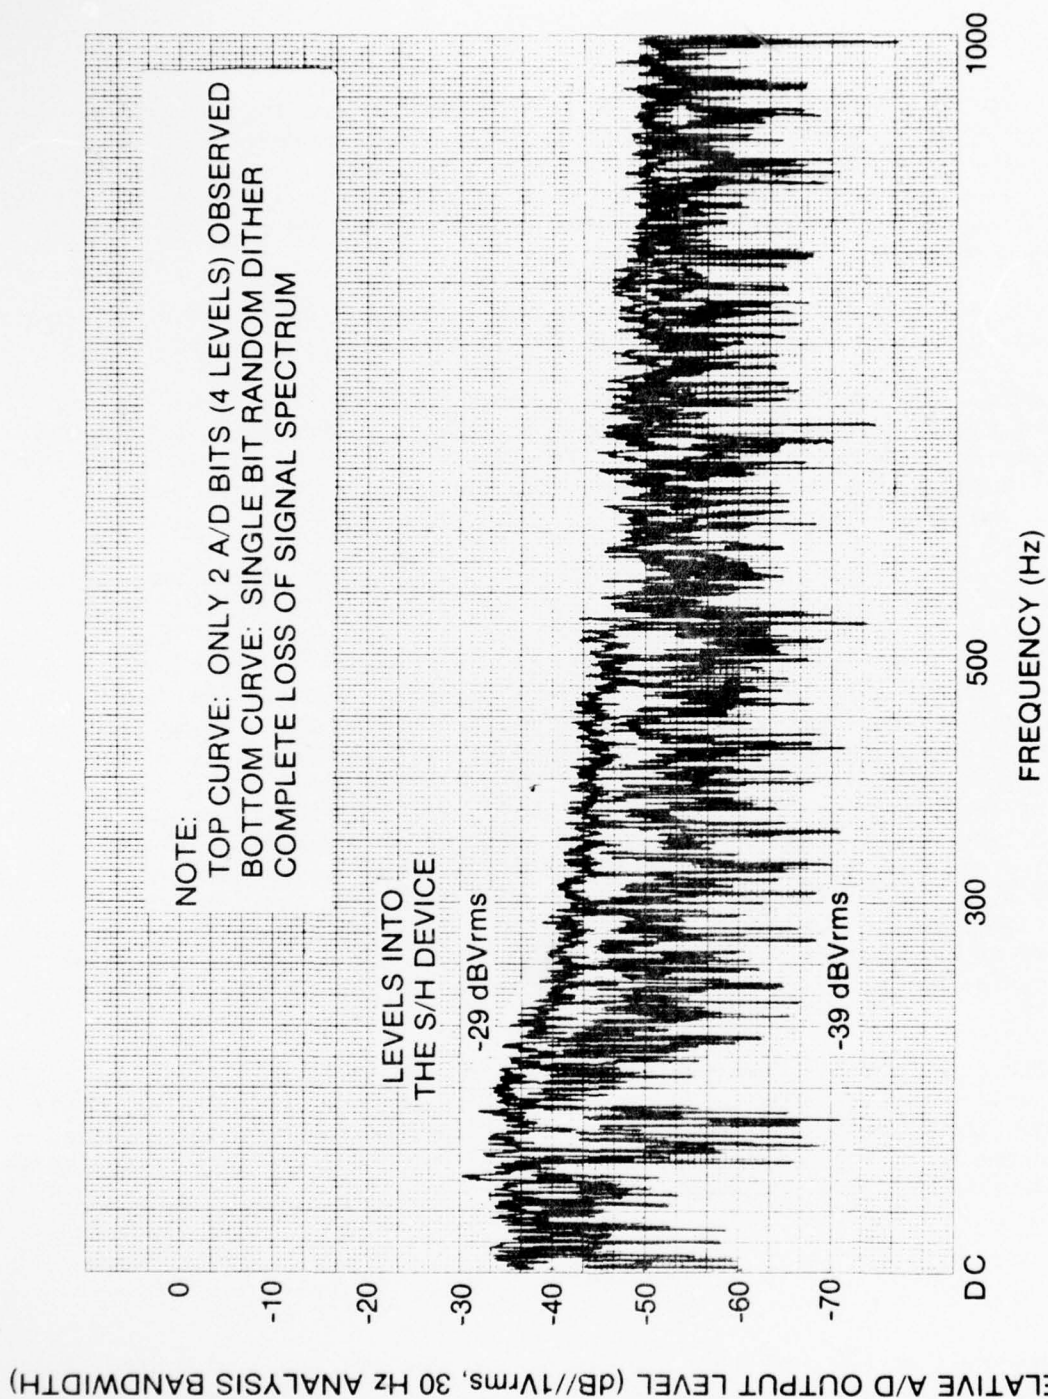
Figure 11C. Quantized Lowpass Filtered Noise, 300 Hz Corner Frequency, 30.0 kHz A/D Sample Rate, Two Input Levels as Indicated

TWO-TONE INPUT

There are several factors to be considered when passing two or more tones through an A/D converter. If we want to obtain minimum spectral distortion, the maximum rms input value that meets this requirement must be established first. For the converter under test here, that value is 4.5 Vrms, or +13 dBV. For K tones at the same spectral level, each tone must be adjusted down by 10 log K to prevent clipping distortion.

Since quantizers are nonlinear devices, multiple tonal inputs are expected to produce sum and difference frequencies (intermodulation) at the output.

At a 3 kHz sample rate, the noise floor was 10 dB higher than in the related figures presented, masking most of the sum, difference, and harmonic components. A sample rate of 30 kHz is used so that these components can be seen. Results at a sampling frequency of 3 kHz may be obtained simply by moving the noise floor up 10 dB.

Two-tone dynamic range is defined as the ability to simultaneously detect large and small amplitude signals. The smallest, single-tone signal which can be detected by this A/D converter was illustrated previously as figure 9 and corresponds to an input level of -40 dBV rms, or 10 mV. The question which arises, however, is whether or not a large signal in the spectral neighborhood of the smaller one will mask the barely detectable signal. The masking could be due to converter desensitizing, generation of nonlinearities, or insufficient analyzer resolution. Figure 12 illustrates two tones, 50 Hz apart, and the resulting spectrum to 1000 Hz. Since +13 dBV rms was established as the optimum, single-tone input level, the two tones were set at 13 dBV - 10 log 2, or 10 dBV rms. It would appear that a two-tone dynamic range of 45 dB (35 dB at a 3 kHz sample rate) is realizable, if a 3 Hz analysis window is used. The presence of the large tone did not increase the noise in the neighborhood of the smaller tone. Except for the harmonic and intermodulation components, the noise floor remained at the same level as for a single tone to 1 kHz.

Figure 13 demonstrates the sensitivity of the difference frequency component to the proximity of the two input tonals. Clearly, as the tonals are separated, the energy in the difference frequency component decreases. It behaves as if the nonlinear process generating the difference frequency is lowpass filtered by a single pole filter at $\omega = 0$.

BAND-LIMITED NOISE

How well an A/D converter reproduces the spectral shape of a band-limited noise source is largely determined by how the A/D converter is biased. A noise signal is not deterministic, and therefore its amplitude must be estimated. Noise standard deviation σ_n may be estimated by using wideband averaging rms meters, such as the HP-403B.
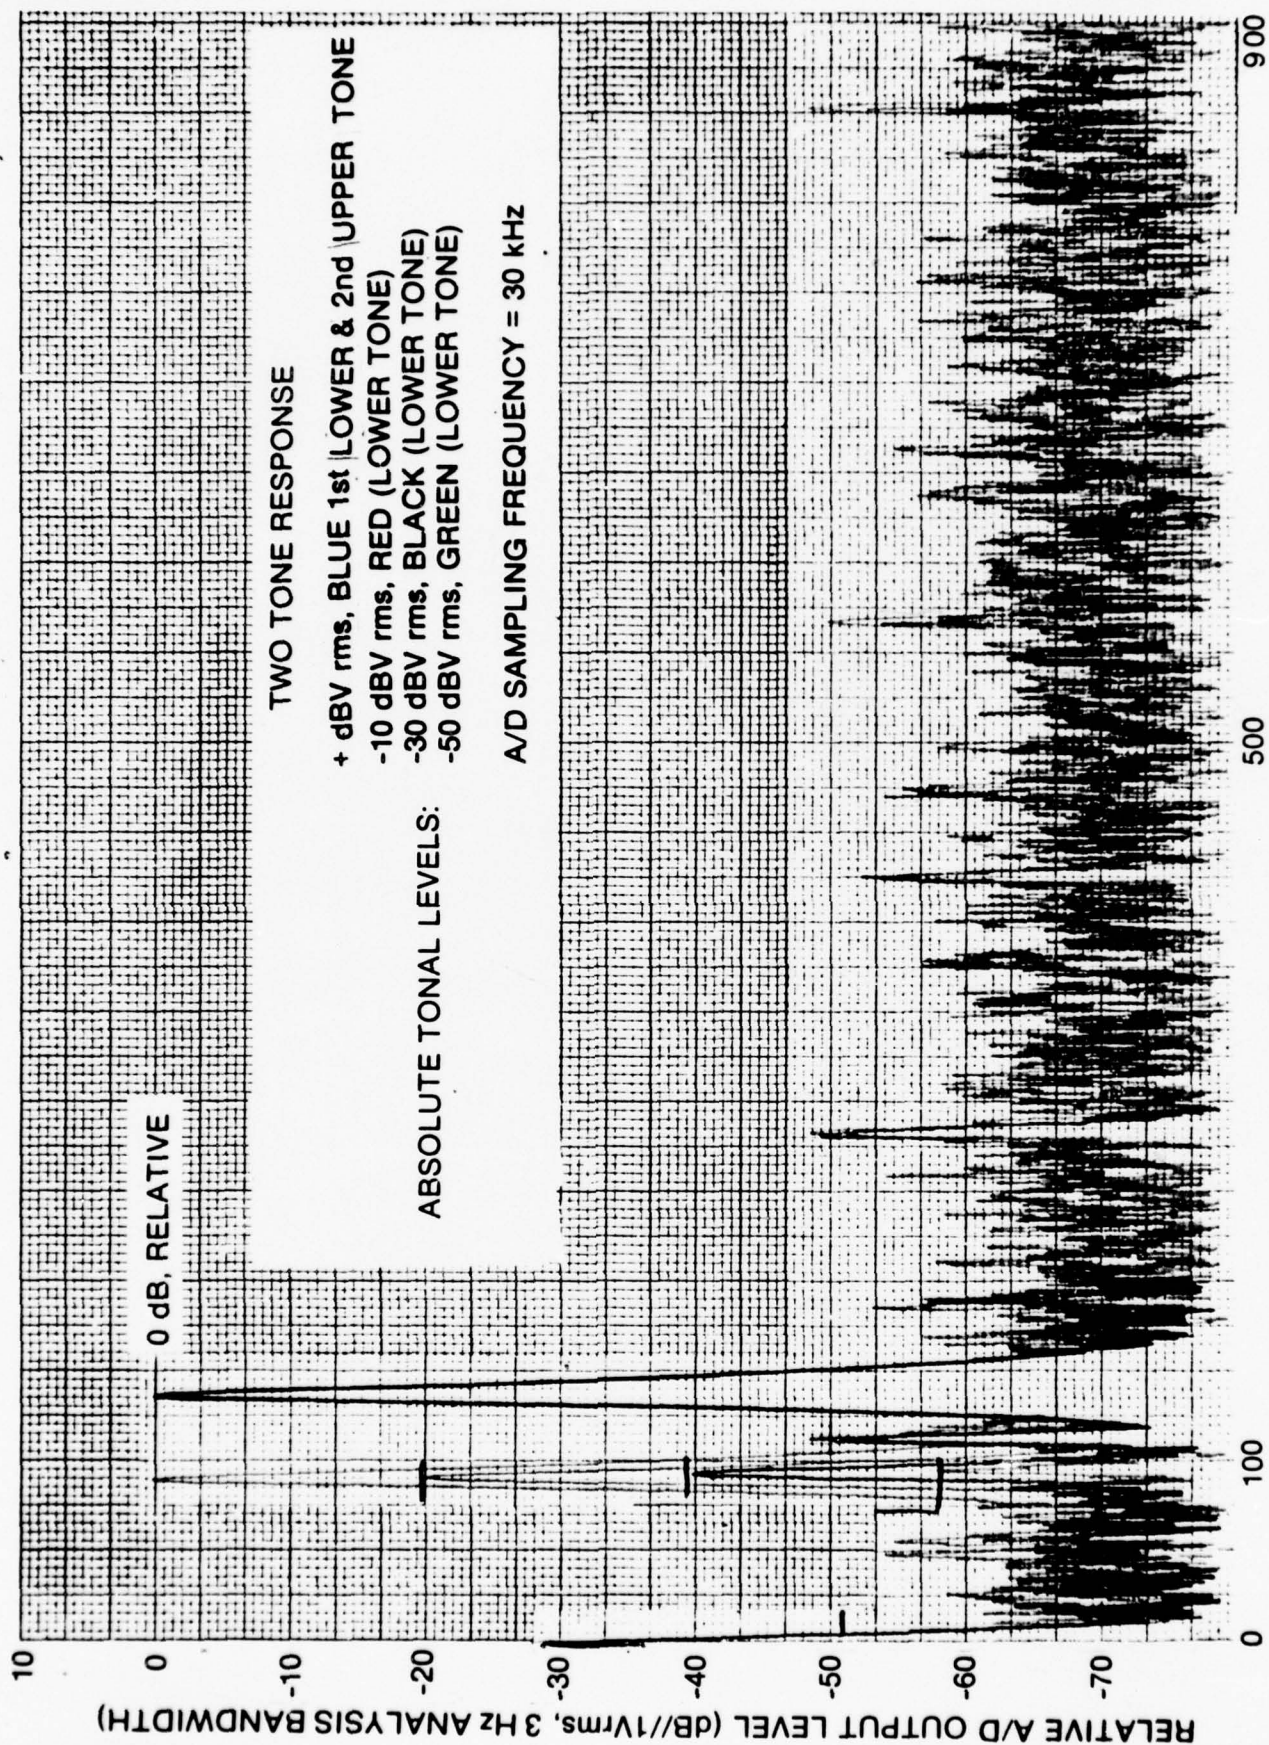
Figure 12. Two-Tone Input, Various Levels

TWO TONE RESPONSE

+ dBV rms, BLUE 1st (LOWER & 2nd (UPPER TONE
-10 dBV rms, RED (LOWER TONE)
-30 dBV rms, BLACK (LOWER TONE)
-50 dBV rms, GREEN (LOWER TONE)

ABSOLUTE TONAL LEVELS:

A/D SAMPLING FREQUENCY = 30 kHz

0 dB, RELATIVE

FREQUENCY (Hz)

RELATIVE A/D OUTPUT LEVEL (dB//1Vrms, 3 Hz ANALYSIS BANDWIDTH)
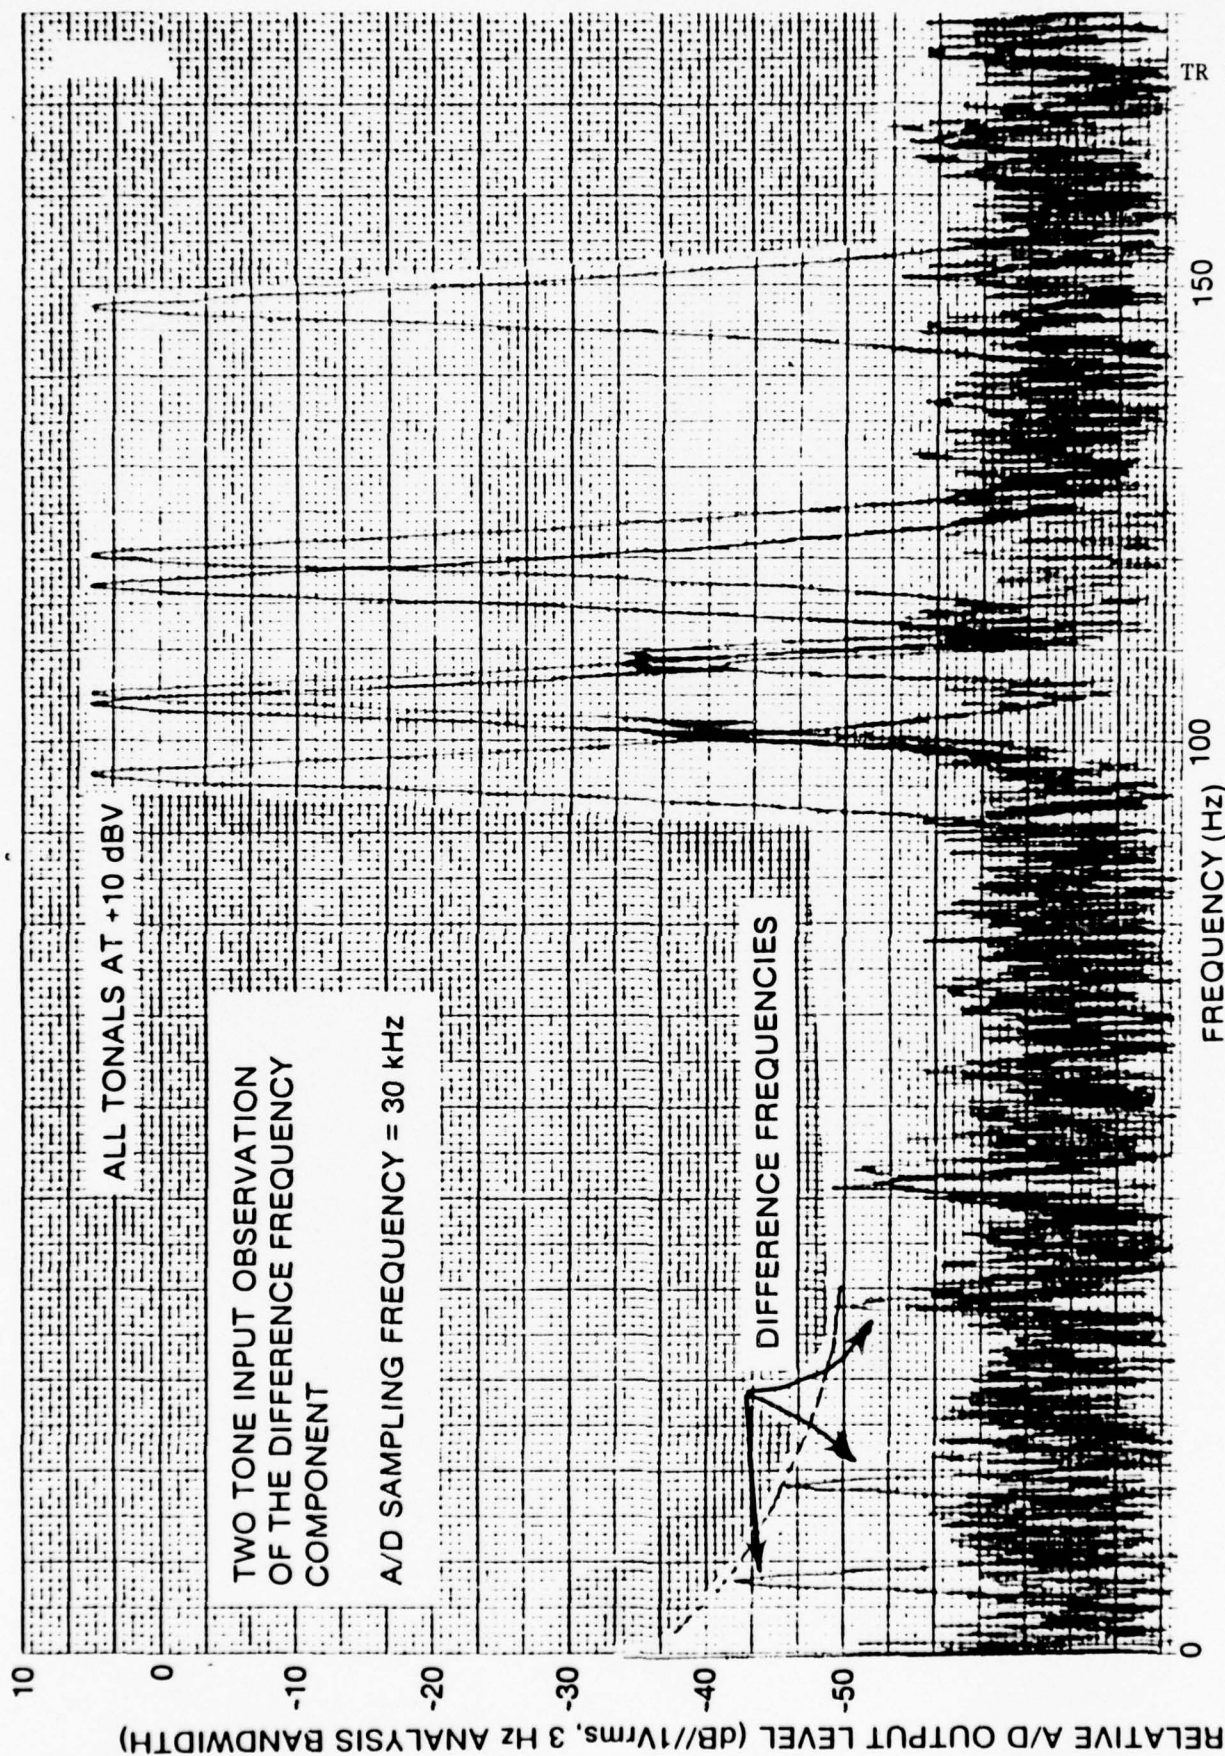
Figure 13. Difference Frequency Component, Two-Tone Input

When dealing with random input signals, it is necessary to examine statistical properties associated with the signal. Expressions which define optimum converter biasing levels that maximize SNR must be developed.

We have already established the quantizer error as

$$\varepsilon_x = \left| \hat{x}(t) - x(t) \right|,$$

where $\hat{x}(t)$ is the quantized version of $x(t)$, a random input signal. The noise power that results is given by

$$P_n = \int_{-\infty}^{+\infty} \left[\hat{x}(t) - x(t) \right]^2 \ p(x) dx \quad ,$$

where the probability density function $p(x)$ is assumed Gaussian (not to be confused with the uniform error distribution in the frequency domain). Another potential source of noise is that due to A/D saturation. The saturation process results in harmonic generation and, therefore, in both a loss of output power at the input signal fundamental frequency and a "filling in" of out-of-band spectral regions that are harmonics or intermodulation components of the noise passband.

These two nonlinearities represent the major sources of A/D quantization noise. A quantitative measure of the power in quantization noise with signal variance, normalized to unity, is shown in figure 14, as a function of the number of bits. The horizontal axis represents the number of standard deviations from the mean saturation level. Saturation is never clearly defined in standard references on this topic; so the results presented here are based upon assumptions described in the section on SIGNAL PLUS LOWPASS FILTERED NOISE. The analytical expressions, which are the basis for these curves, are derived by Bennett.[2]

Note that the expression for P_n is actually a calculation of the quantization error variance, such that the curves of figure 15 are actually $\sigma_\varepsilon^2/\sigma_x^2$ or 1/SNR for a normalized input. The calculated value for P_n at the lowest possible point on the 6-bit quantizer curve is $20 \log (1 \times 10^{-3})$ or -30 dB from dc to the Nyquist frequency. It is common practice to bias the A/D converter input to the right of the low point on the curves, since saturation noise degrades SNR much more rapidly than quantization noise.

Figure 11 shows four curves of band-limited noise spectra at various input levels. Figure 15 is the band-limited noise spectra before quantizing. Should desired tonals reside on the noise slope, any decrease in the roll-off slope due to saturation "fill-in" may mask the signal.

The top curve in figure 11A represents a band-limited noise input signal adjusted to +13 dBV rms on a broadband rms meter. This is at a level just over saturation. At +13 dBV, clipping was observed on the oscilloscope "very infrequently." This is hardly a scientific measurement, but it was not possible to record the frequency of clipping occurrence. If 13 dBV (4.4 Vrms) is defined as the mean saturation level, to be on the safe side, this corres-
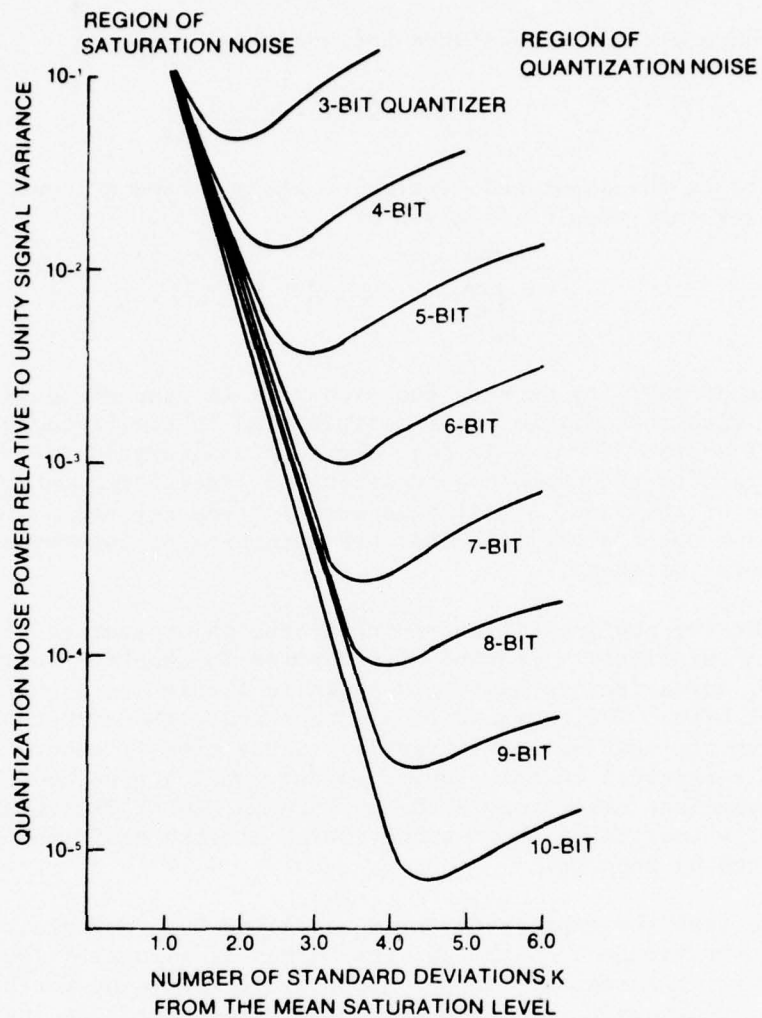
25

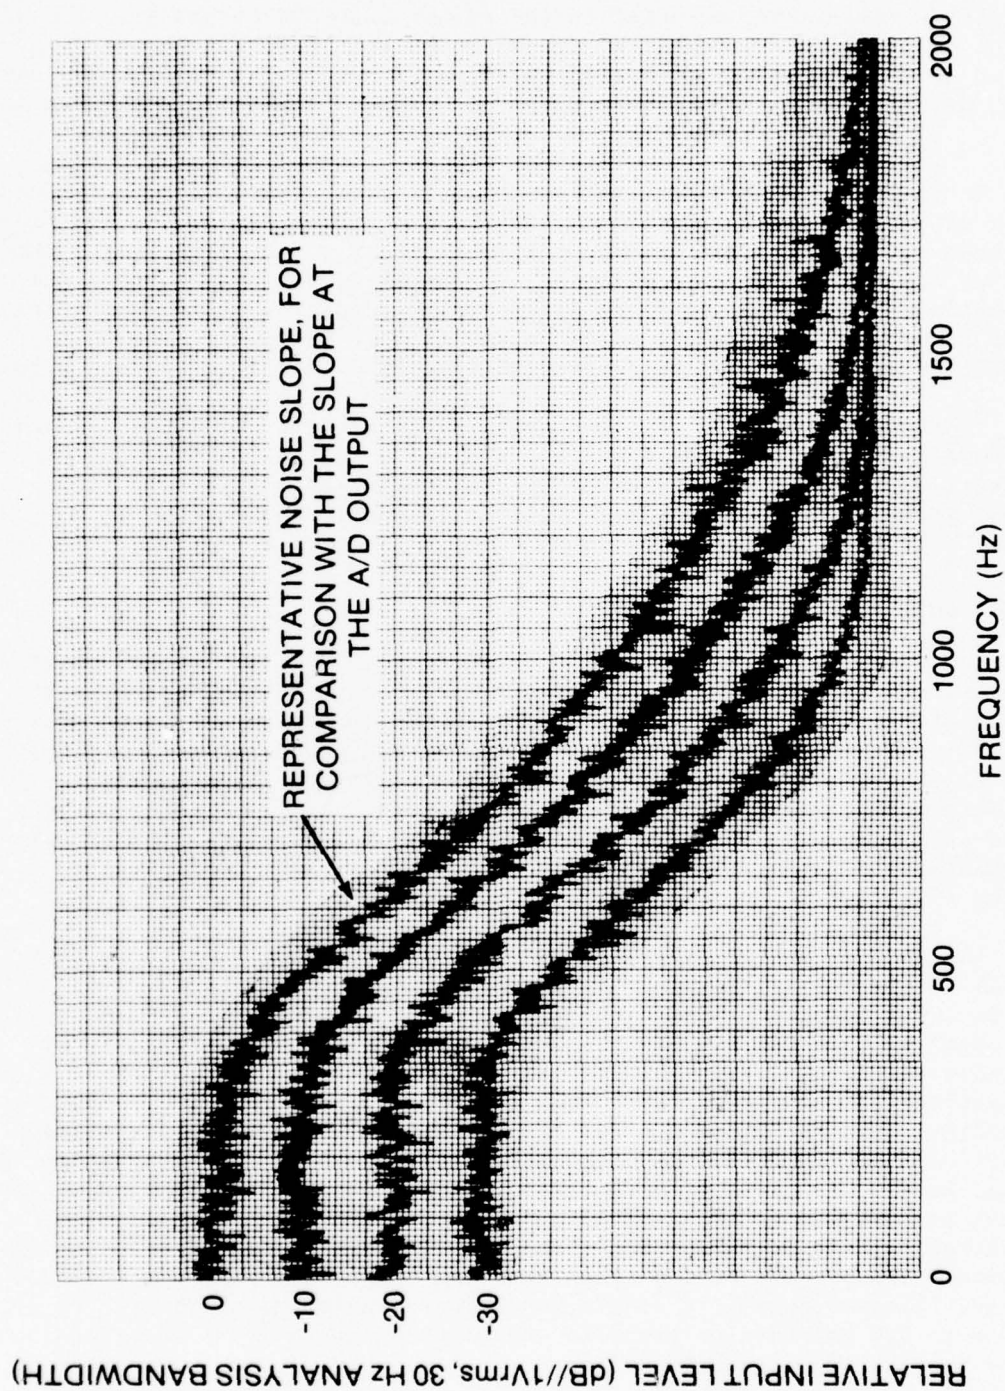Figure 14.   Power in the Error Due to Quantization and Saturation

Figure 15. Lowpass Filtered Noise, Input Spectra, Four Levels

ponds to K = 4.4/3.55 = 1.24. A comparison of the noise slopes in figure 14
to the quantized spectrum in figure 11A demonstrates that a good deal of
saturation fill-in has occurred in the region from 300 to 600 Hz. The effect
is not nearly so catastrophic as is predicted by figure 15, however, perhaps
due to the conservative definition of the mean saturation level. The next
curve down accurately reproduces the noise slope. Its input rms level is +1
dBV.

In figure 11B, the quantization noise in the stopband for the second
curve is lower than for the other cases, which indicates that dynamic range
has been optimized. The increase of the stopband noise level in the remain-
ing two cases indicates operation of the converter too far to the right of
the optimum bias level, and a corresponding increase in quantization noise
power is noted. This condition is not observable in figure 11A, where the
Nyquist frequency is less than three octaves from the filter cutoff frequency.

In the optimum case, the level difference between noise passband and
stopband is about 30 dB in figure 11A and 40 dB in 11B. There are not bandwidth
corrections to be made here, since we are looking at the ratio of two noise
levels derived from an identical analyzer resolution bandwidth.

Signal Plus Lowpass Filtered Noise

Two test configurations were considered. In the first test, noise was
lowpass filtered with the corner frequency set at 200 Hz. The noise was
summed with a tonal at 275 Hz, which is in the transitional region of the
filter. The SNR at the filter output was fixed, in the spectral neighbor-
hood of the signal, to +30 dB. Refer to test configuration 1, figure 16.
The second test was similar to the first, except that a second tonal was
added at 130 Hz. Both tonals were set such that the SNR, in the spectral
neighborhood of each signal, was +10 dB. Refer to test configuration 2,
figure 17.

The purpose of the first test was to: (1) demonstrate the sensitivity
of SNR to biasing at the A/D converter input, and (2) plot experimental SNR
versus level input for comparison with Bennett's calculations.[2] As previously
discussed in this report, power in the quantization noise/unity signal
variance was identified as a normalized curve of $(SNR)^{-1}$ vs K standard
derivations from the mean saturation level. It was desired to observe a
correlation between the shapes of the analytically obtained, normalized
$(SNR)^{-1}$ and the experimental data. Any experimental SNR data in which the
signal levels are fixed, and the noise in the spectral neighborhood of the
signal is subject to the previously observed effects of saturation and
quantization, should follow the $(SNR)^{-1}$ curves (figure 15). Examples of this
condition are signals on the slope near the stopband, or in the stopband, of
lowpass filtered noise. A single tonal in the noise passband will be relatively
immune to SNR degradation provided the passband noise level is maintained
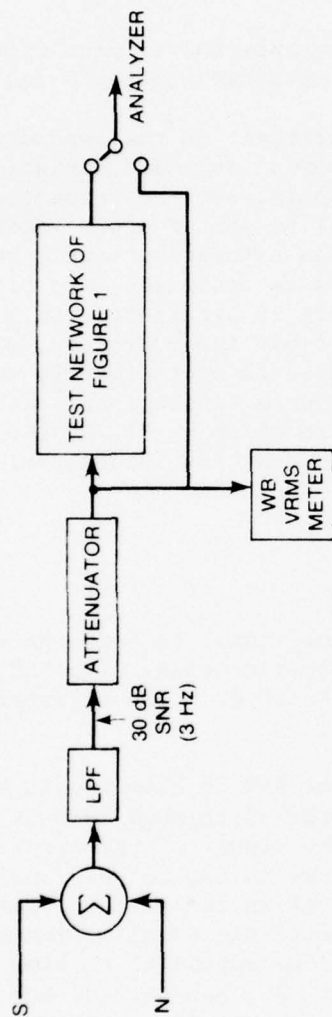above the quantization noise floor.
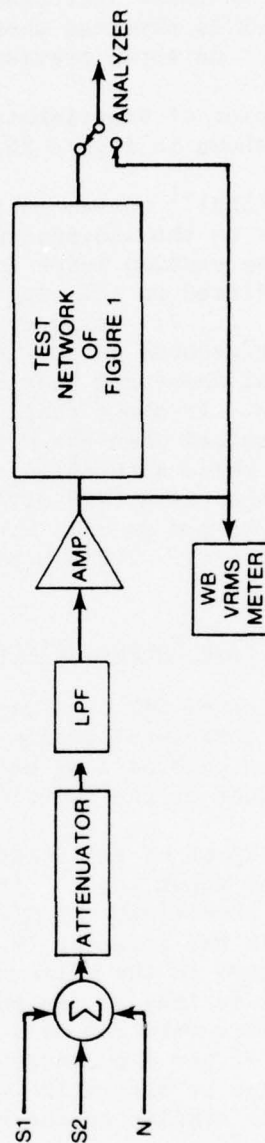
28

Figure 16. Test Configuration 1

Figure 17. Test Configuration 2

29

Figures 18B through 23 illustrate the spectrum in figure 18A after it is passed through the A/D converter. For high input levels, transitional and stopband regions fill in with saturation noise, degrading the SNR. At low input levels, the shape of the lowpass filtered noise disappears as the quantization noise increases relative to the signal, again degrading the SNR. Optimum SNR is obtained when the noise level is adjusted between 0 and +5 dB broadband. As shown previously, this level best reproduces the noise slope.

The plot of experimentally obtained SNR versus noise (plus signal) input level is shown in figure 24. Optimum SNR is obtained between 0 and +5dB

The $(SNR)^{-1}$ curves in figure 15 are logarithmic on the dependent axis but linear on the independent axis. The number of standard deviations from the mean saturation level is measured in decibels; so it is expected that $K\sigma$ (when fitted to the experimental data) will be spaced apart logarithmically, $K = 1, 2, \ldots, 7$. It is evident that the "mean saturation level" referred to by many references on this subject is intended to designate hard clipping. The general shape and peak of the SNR curve are in agreement with analytical prediction. It shows that, in the case of a 6-bit A/D converter, maximum SNR is obtained when the noise level is adjusted to approximately one-quarter the value where saturation distortion begins for a single tonal. The 2σ point on the curve (+18 dBV rms) corresponds to clipping 31.7% of the time, 3σ to 0.27%, and so on. Note that these broadband, rms input levels are below 14 V peak. The A/D power supply is $\pm$ 15 Vdc.

Two Tones Plus Lowpass Filtered Noise

The second S+N test was performed with two tones, in both the passband and the transitional region of the lowpass filtered noise. The SNR, in the spectral neighborhood of each tone, was set to +10 dB. The difference between the amplitude of the two tones was 20 dB.

The object of this test was to observe the SNR of both tonals while varying the input level. The results in figures 25 through 29 show an increased sensitivity to A/D bias level for the signal of the transitional region. In the figures, TR refers to the signal on the noise slope and PA to the signal in the noise passband. The signal in the passband remains at +10 dB SNR as long as the passband level is above the quantization noise floor. Further tests were run on a hard clipper (see the appendix) to simulate operation of the A/D converter in full saturation. Behavior of the noise spectrum due to saturation and quantization noise, as illustrated in the figures, is similar to the results obtained previously.
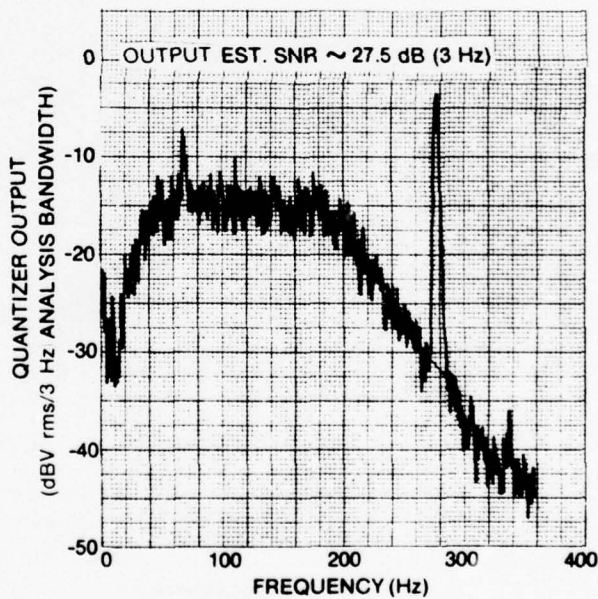
A. Spectrum at the Quantizer Input
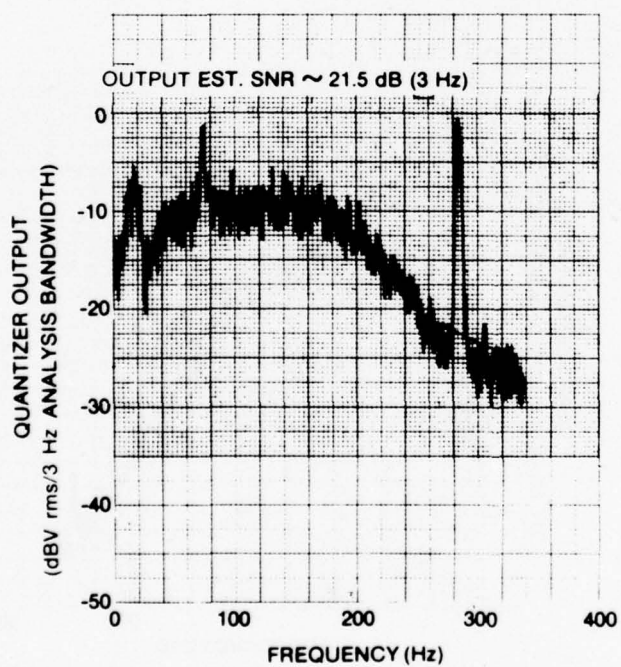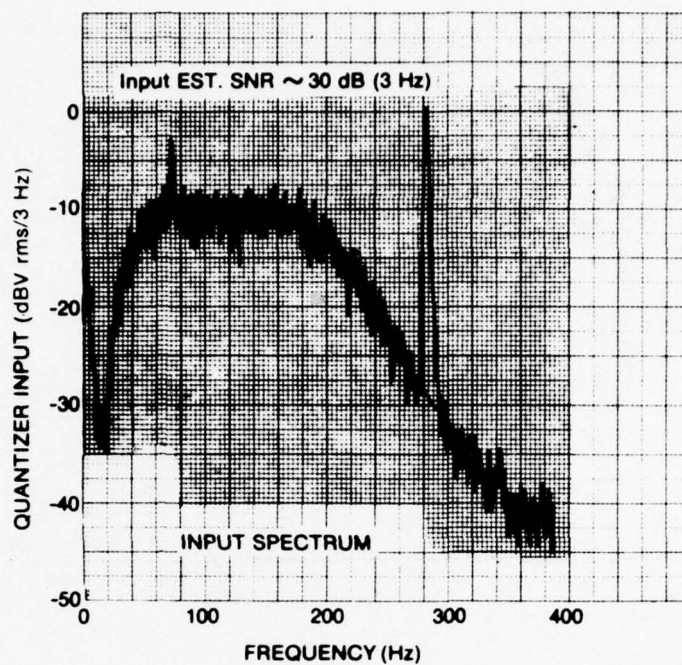
B. Noise Level +17 dBV rms, Signal Level +7 dBV rms

Figure 18. SNR Measurement: Signal in Transitional Region, Lowpass Filter Set at 200 Hz, Signal at 275 Hz; Reference Test Configuration 1
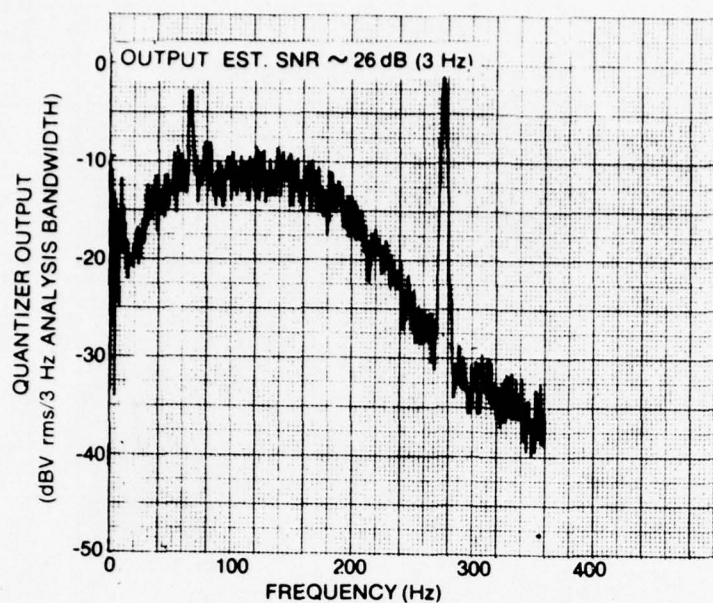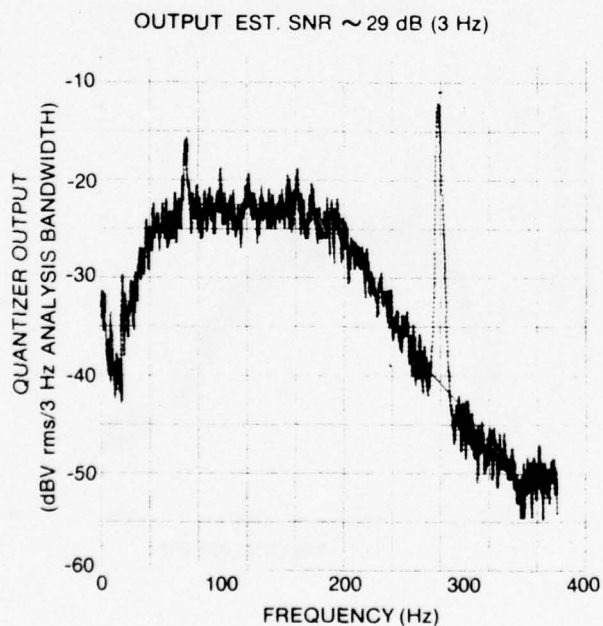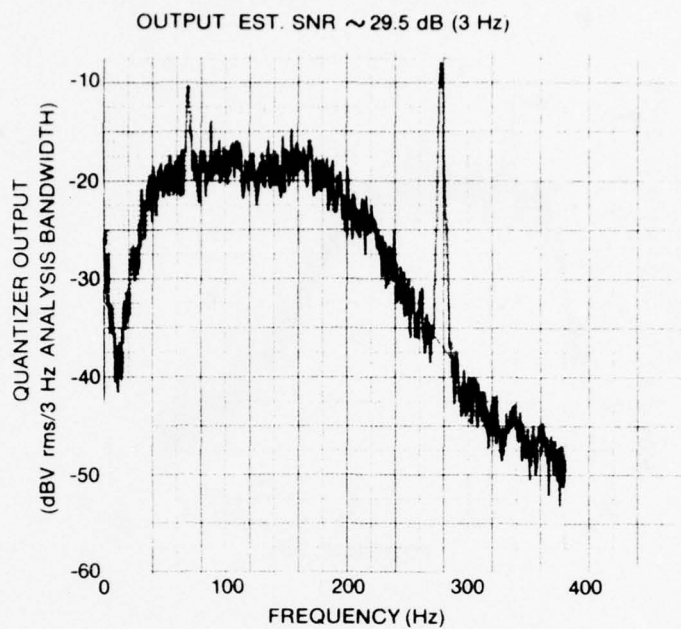
A. Noise Level at
+15 dBV rms, Signal
Level at +5 dBV rms

B. Noise Level at
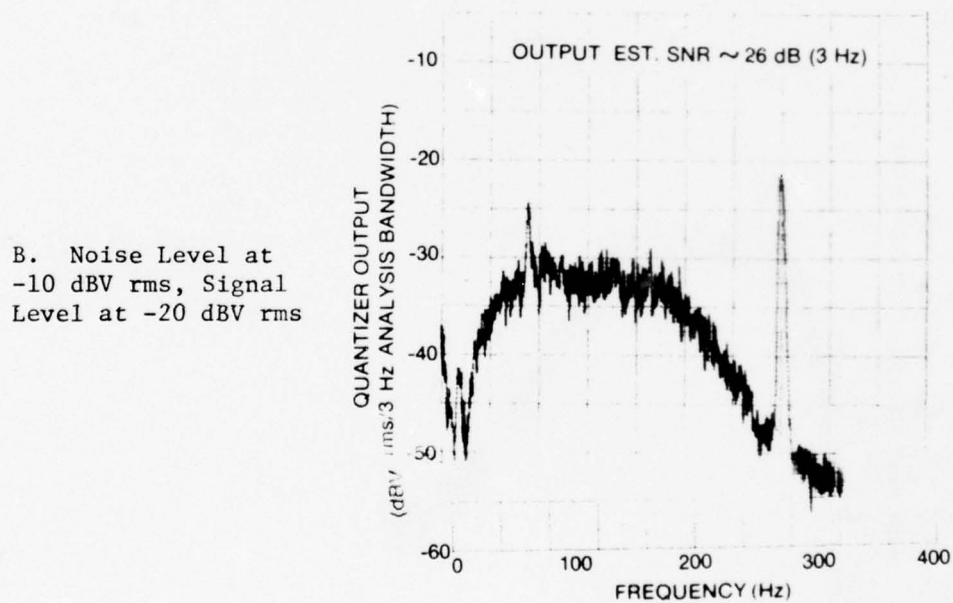+10 dBV rms,
Signal Level at 0 dBV rms

Figure 19. SNR Measurement: Signal in Transitional Region,
Lowpass Filter Set at 200 Hz, Signal at 275 Hz;
Reference Test Configuration 1

OUTPUT EST. SNR ~ 29.5 dB (3 Hz)

A. Noise Level at +5 dBV rms, Signal Level at -5 dBV rms

OUTPUT EST. SNR ~ 29 dB (3 Hz)

B. Noise Level at 0 dBV rms, Signal Level at -10 dBV rms

Figure 20. SNR Measurement: Signal in Transitional Region,
Lowpass Filter Set at 200 Hz, Signal at 275 Hz;
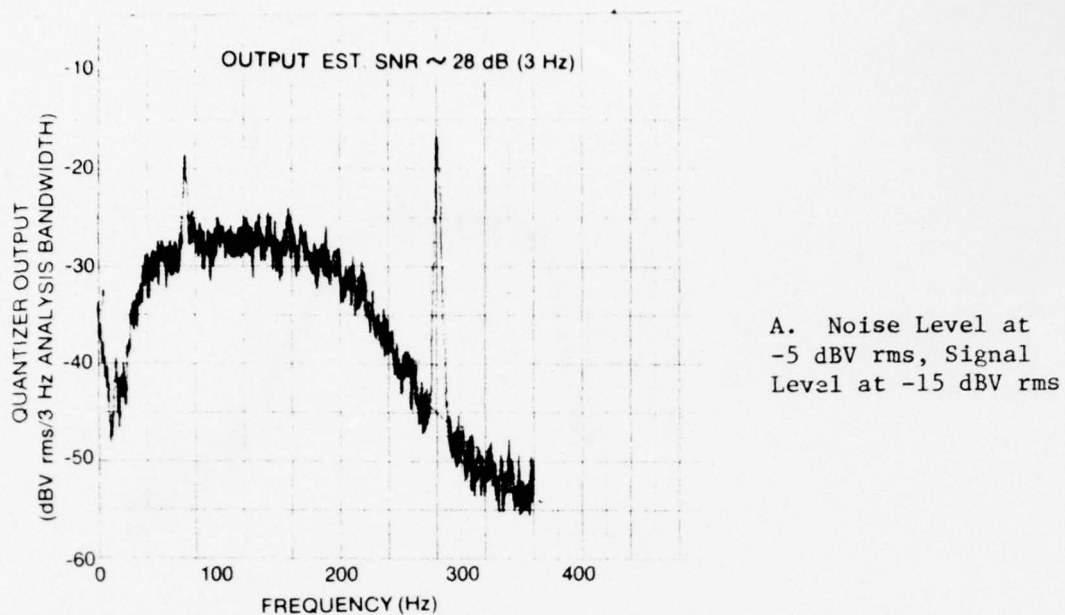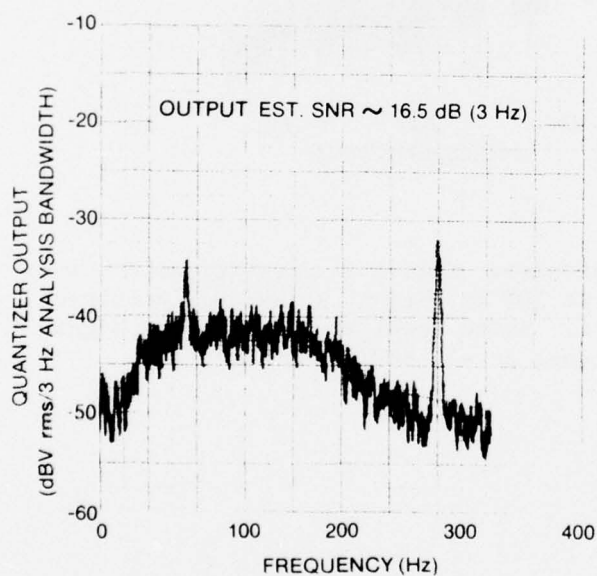Reference Test Configuration 1

33

OUTPUT EST. SNR ~ 28 dB (3 Hz)

A. Noise Level at
-5 dBV rms, Signal
Level at -15 dBV rms

OUTPUT EST. SNR ~ 26 dB (3 Hz)

B. Noise Level at
-10 dBV rms, Signal
Level at -20 dBV rms

Figure 21. SNR Measurement: Signal in Transitional Region,
Lowpass Filter Set at 200 Hz, Signal at 275 Hz;
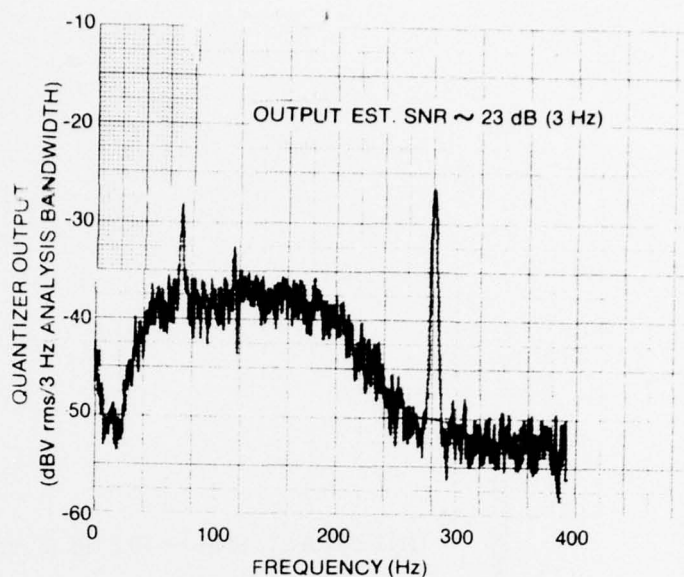Reference Test Configuration 1

34

A. Noise Level at
-15 dBV rms, Signal
Level at -25 dBV rms

B. Noise Level at
-20 dBV rms, Signal
Level at -30 dBV rms

Figure 22. SNR Measurement: Signal in Transitional Region,
Lowpass Filter Set at 200 Hz, Signal at 275 Hz;
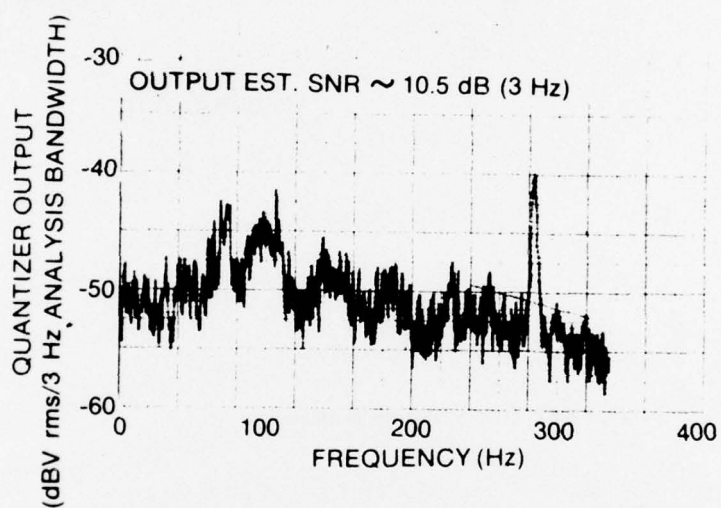Reference Test Configuration 1

Figure 23. SNR Measurement: Signal in Transitional Region, Lowpass Filter Set at 200 Hz, Signal at 275 Hz; Reference Test Configuration 1. Noise Level at -25 dBV rms, Signal Level at -35 dBV rms
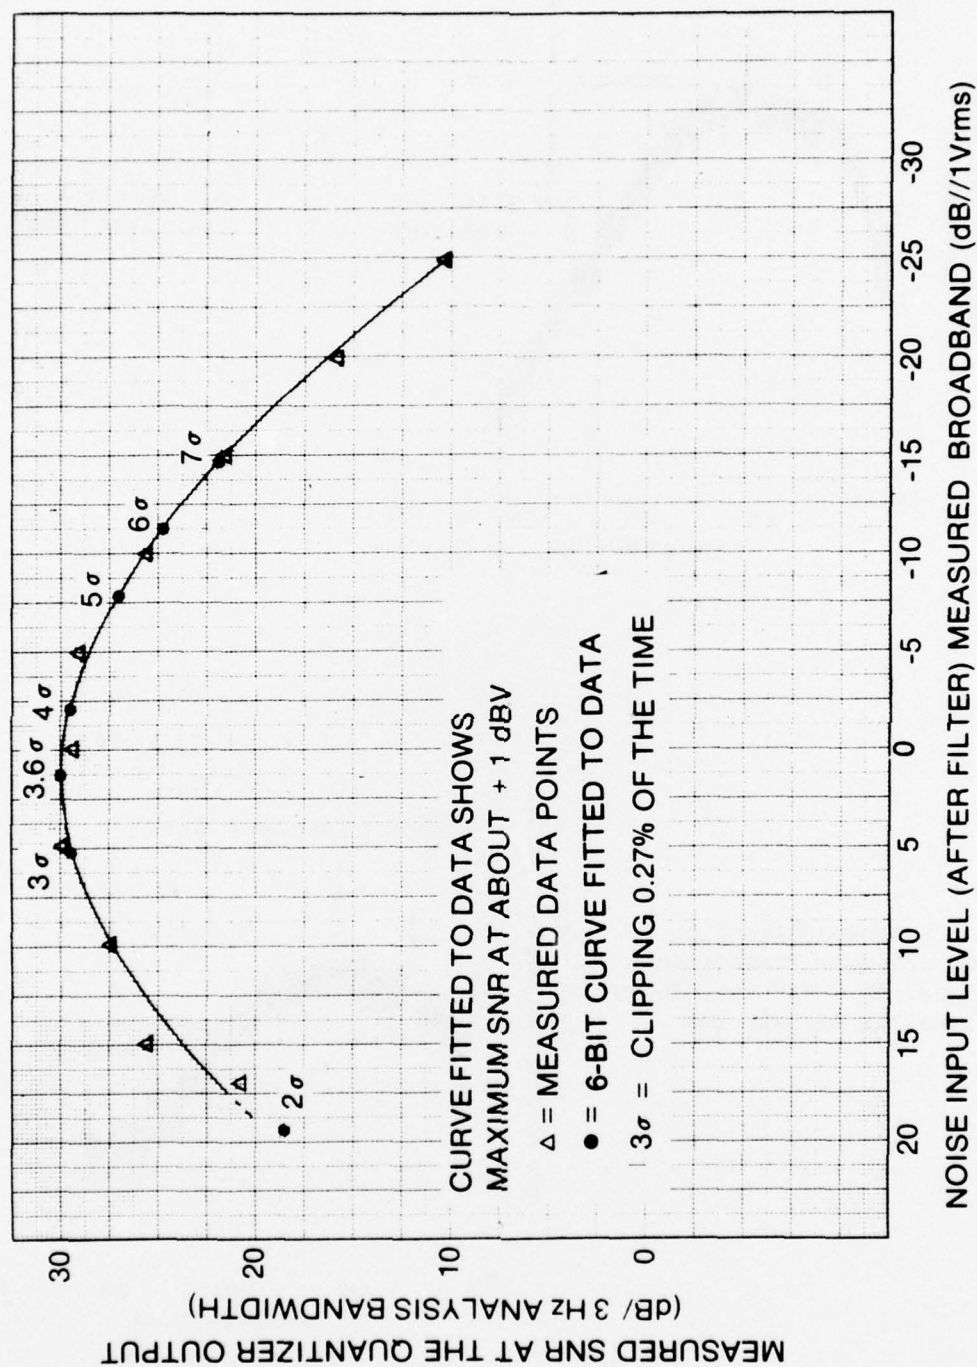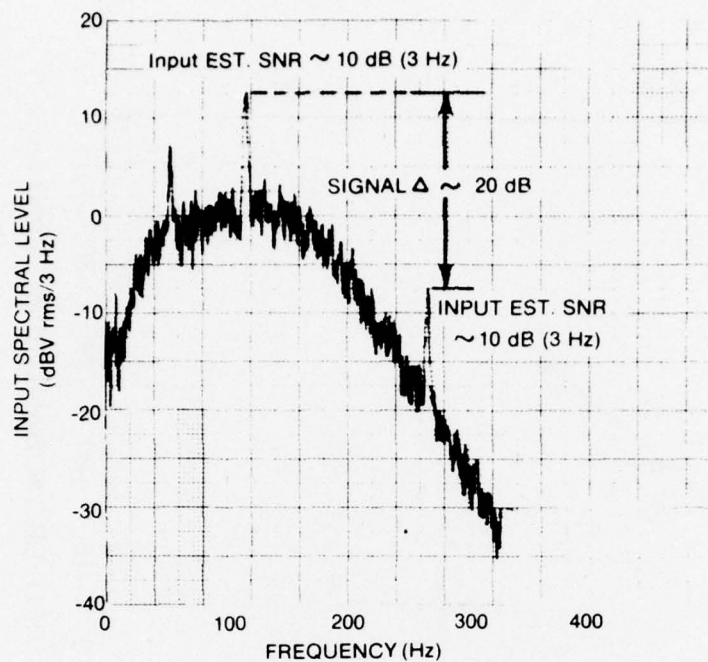
Figure 24. Measured SNR at the Quantizer Output With
The 6-Bit Curve From Figure 15 Fitted to Data

A. Spectrum at the Quantizer Input

B. Passband Signal at +15 dBV rms, Transitional Signal at 0 dBV rms, Noise Signal at +17 dBV rms
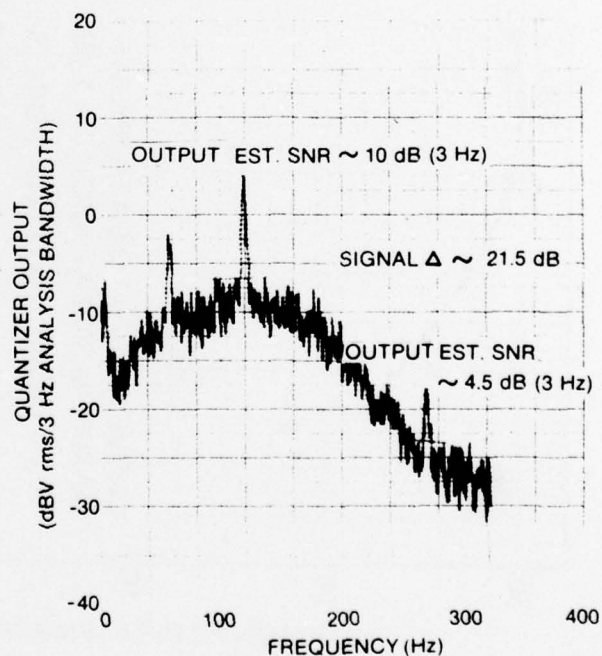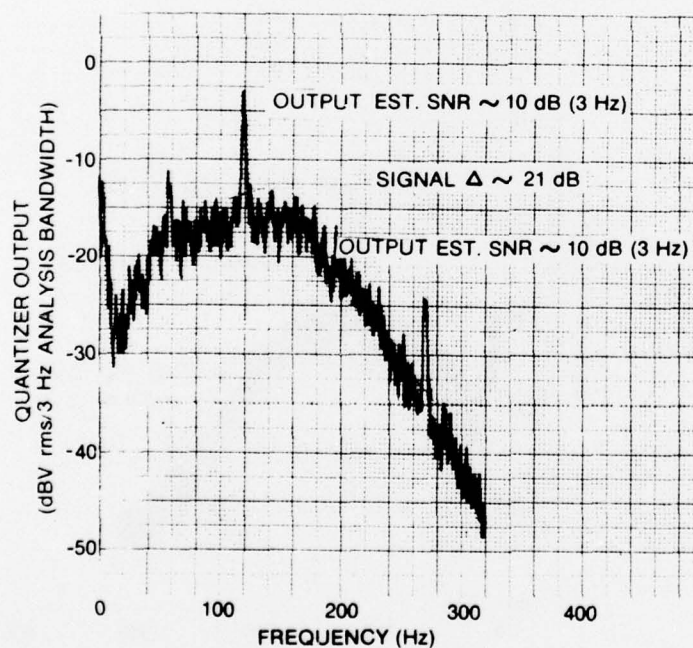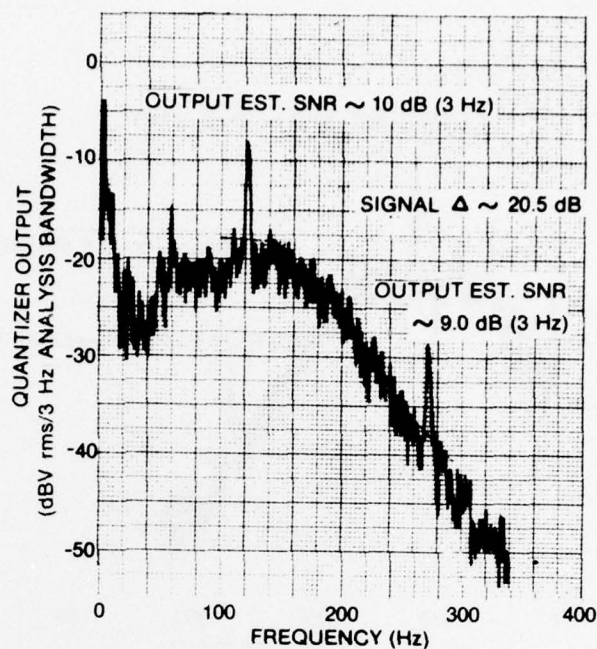
Figure 25. SNR Measurement, Two Tones Plus Lowpass Filtered Noise, LPF Noise at 200 Hz; Passband Signal at 130 Hz, Transitional Band Signal at 280 Hz; Reference Test Configuration 2

A. Passband Signal at +5 dBV rms, Transitional Signal at −10 dBV rms, Noise Signal at +7 dBV rms

B. Passband Signal at 0 dBV rms, Transitional Signal at −15 dBV rms, Noise Signal at +2 dBV rms

Figure 26.   SNR Measurement, Two Tones Plus Lowpass Filtered Noise, LPF Noise at 200 Hz; Passband Signal at 130 Hz, Transitional Band Signal at 280 Hz; Reference Test Configuration 2

39

A. Passband Signal at
-5 dBV rms, Transitional
Signal at -20 dBV rms,
Noise Signal at -3 dBV rms

B. Passband Signal at
-10 dBV rms, Transitional
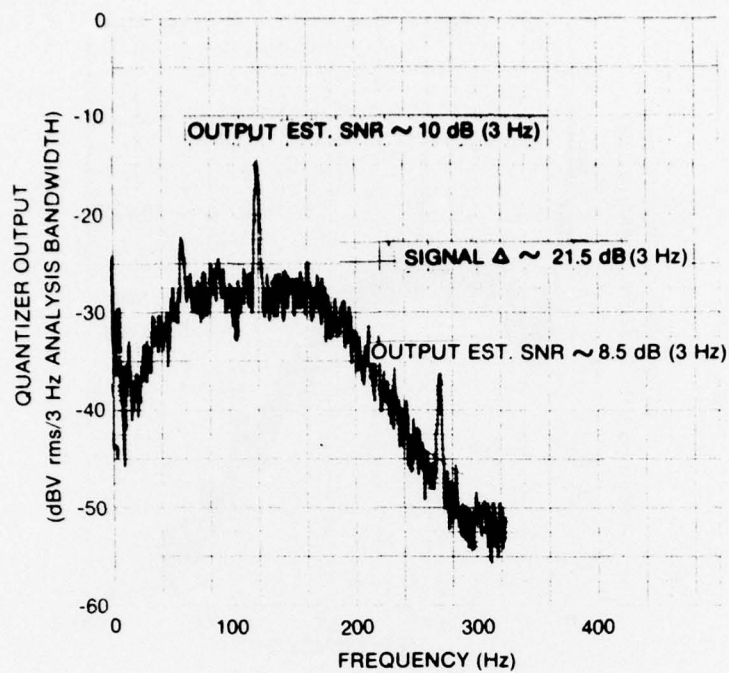Signal at -25 dBV rms,
Noise Signal at -8 dBV rms

Figure 27. SNR Measurement, Two Tones Plus Lowpass Filtered Noise,
LPF Noise at 200 Hz; Passband Signal at 130 Hz, Transitional
Band Signal at 280 Hz; Reference Test Configuration 2

A. Passband Signal at −15 dBV rms, Transitional Signal at −30 dBV rms, Noise Signal at −13 dBV rms
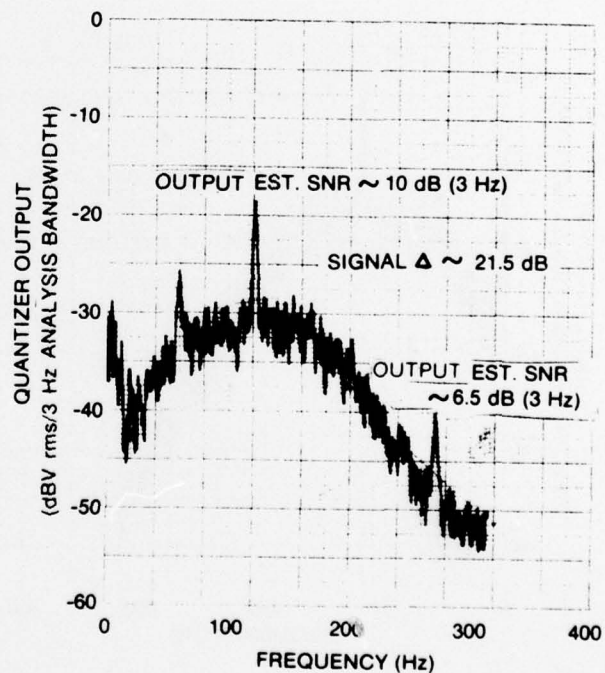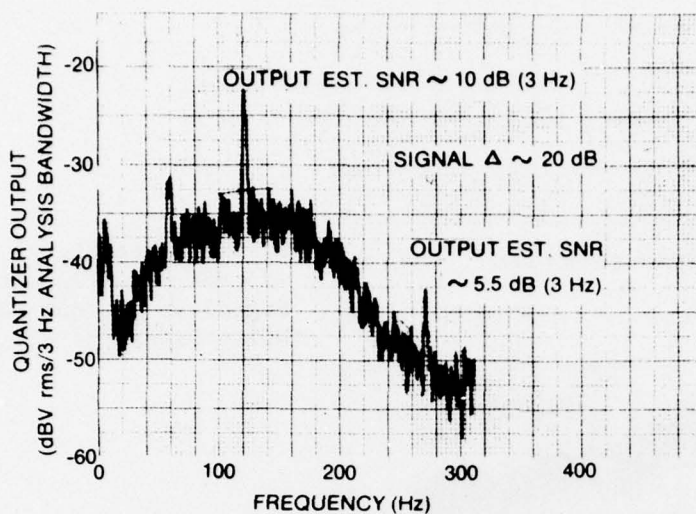
B. Passband Signal at −20 dBV rms, Transitional Signal at −35 dBV rms, Noise Signal at −18 dBV rms

Figure 28. SNR Measurement, Two Tones Plus Lowpass Filtered Noise, LPF Noise at 200 Hz; Passband Signal at 130 Hz, Transitional Band Signal at 280 Hz; Reference Test Configuration 2

A. Passband Signal at
−25 dBV rms, Transitional
Signal at −40 dBV rms,
Noise Signal at −23 dBV rms

B. Passband Signal at
−30 dBV rms, Transitional
Signal at −45 dBV rms,
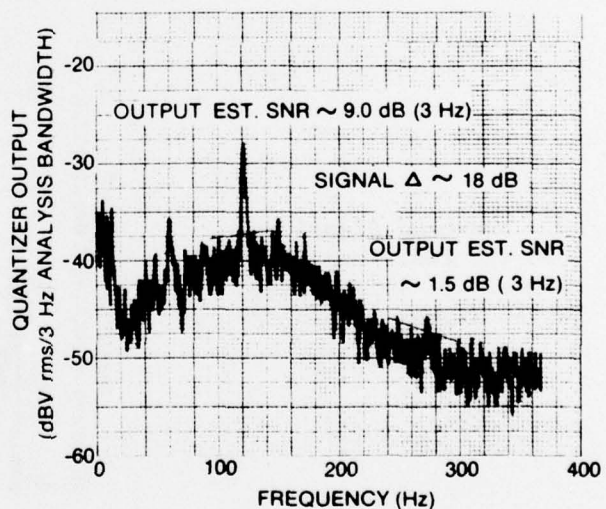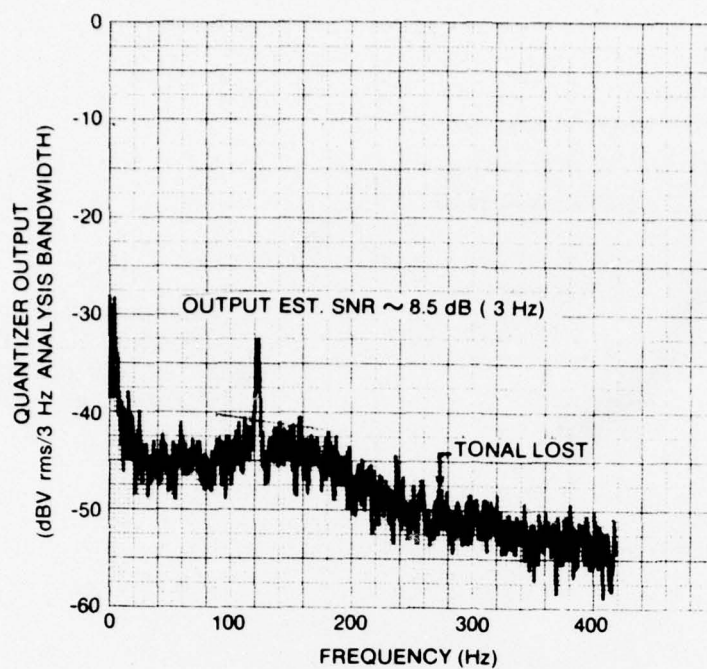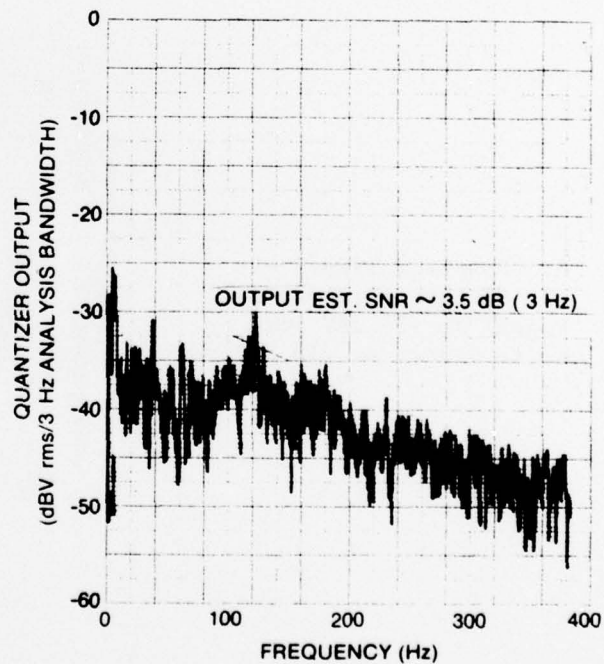Noise Signal at −28 dBV rms

Figure 29. SNR Measurement, Two Tones Plus Lowpass Filtered Noise,
LPF Noise at 200 Hz; Passband Signal at 130 Hz, Transitional
Band Signal at 280 Hz; Reference Test Configuration 2

CONCLUSIONS

- At a 3.0 kHz sampling frequency, the 5-bit-plus-sign A/D converter demonstrated a full tone signal-to-noise ratio (SNR) of about 69 dB in a 1 Hz band.

- Two-tone dynamic range is shown to be about 35 dB in a 3 Hz analysis window. This figure also represents the maximum relative difference obtainable between bandpass filtered noise and stopband levels.

- In applications where the system passband excludes harmonic frequencies (one octave or less), an A/D converter can be driven slightly into saturation and an increase in SNR is realized at the sacrifice of absolute levels.

- Analytical predictions of SNR for a single tone appeared to be about 6 dB less than the measured SNR, due to the ambiguity of a large noise variance. The predicted SNR is shown to be accurate, however, by analyzing the case of a band-limited noise input and by adjusting the appropriate levels. The spectra of quantized tonals appear to have an average noise floor lower than predicted due to the many discrete noise spikes that result. Refer to figures 2 through 9.

- Given a signal in a band-shaped (lowpass filtered) noise environment, maximum SNR is realized when the spectral shape of the noise is most accurately reproduced. In the case of the 5-bit-plus-sign A/D converter analyzed here, that condition is achieved when the rms noise is set to one-quarter of the level that causes a sharp increase in harmonic energy for a pure sinusoid input. Above this level, saturation causes a filling-in of the spectral regions that are harmonics or intermodulation components of the noise passband. Below this level, quantization noise is shown to increase. These results are predicted in the work of Bennett,[2] and others.[3,4]

- The analytical quantization noise predictions by Bennett are demonstrated to fit experimental SNR data for a signal in the transitional region of lowpass filtered noise.

- The dynamic range of an A/D converter with only the LSB activated is shown to be about 11 dB. This result is also observable in the familiar analytical expression (which assumes a Gaussian input) for quantizer mean square error,

$$\sigma_\epsilon^2 = \frac{2^{-2b}}{12} \ ,$$

or

$$\text{Noise power} = 10 \log 2^{-2b} - 10 \log 12,$$

so noise power $\cong (-6b - 10.8)$dB. This expression predicts noise power 10.8 dB below the LSB.

• The 5-bit-plus-sign A/D converter discussed here was implemented by truncating the 6 low order bits of an 11-bit-plus-sign, offset binary A/D converter. Zero to positive or negative to positive transitions with 1/64 the magnitude of the redefined LSB would cause the sign bit to dither, as would a hard clipper. This extended the spectral dynamic range beyond the predicted values, since sign information was preserved.

APPENDIX

BAND-SHAPED NOISE SPECTRA THROUGH A HARD CLIPPER

Additional tests were performed on a hard clipper to observe operation of an A/D converter in hard saturation, without the second-order effects due to overloading. The clipper also demonstrates A/D operation at low levels, which causes dithering of the sign or of the LSB, depending on converter type.[3,4] The signal phase is preserved, and meaningful spectral information may be extracted with the limited dynamic range of a clipper. It is interesting to note that the operation of a hard clipper closely approximates n-bit A/D converter operation at both hard saturation and below LSB levels.

Observation of the band-shaped noise spectra through a clipper (figure A-1) presents the 11 dB dynamic range obtainable. A parallel can be drawn to the analytical prediction of quantization noise variance[1]:

$$\sigma^2_{\varepsilon_x} = \int_{-\frac{1}{2}lsb}^{+\frac{1}{2}lsb} \varepsilon^2_x \, d\varepsilon_x = \frac{2^{-2b}}{12}$$

so that

$$\sigma_{\varepsilon_x} = \frac{2^{-b}}{\sqrt{12}} \quad .$$

The quantization noise predicted by this is

$$P_{\varepsilon_x} = 10 \log \frac{2^{-b}}{\sqrt{12}} = 10 \log 2^{-b} - 10.8 \text{ dB (rms)}.$$

If it is noted that 2^{-b} is actually the smallest converter step size, where full scale is normalized to unity, it can be seen that the predicted noise floor is about 11 dB below the LSB.
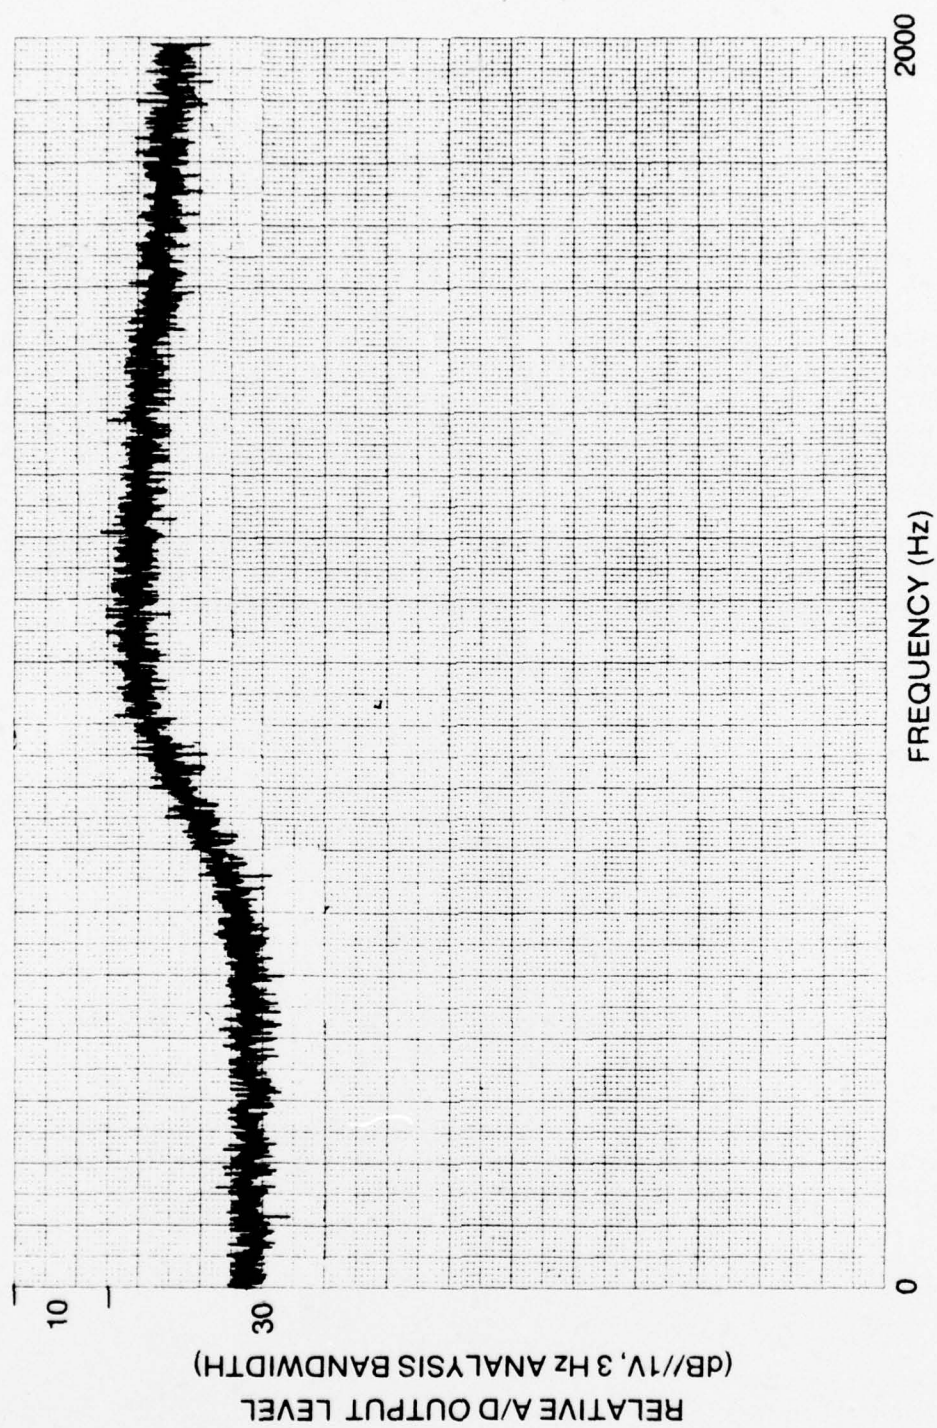
Figure A-1. Band-Shaped Noise Spectrum Through a Hard Clipper

REFERENCES

1. R. K. Otnes and L. Enochson, Digital Time Series Analysis, Wiley, New York, 1972.

2. W. R. Bennett, "Spectra of Quantized Signals," Bell System Technical Journal, vol. 27, July 1948, pp. 446-472.

3. G. A. Gray and G. W. Zeoli, "Quantization and Saturation Noise Due To Analog-to-Digital Conversion," IEEE Transactions on Aerospace And Electronic Systems, vol. AES-7, no. 1, January 1971, p. 222.

4. J. Max, "Quantizing for Minimum Distortion," IRE Transactions on Information Theory, vol. IT-6, no. 1, March 1960, pp. 7-12.

5. A. H. Nuttall, Waveforms and Spectrum Distortion of Narrowband Signals in Noise Caused by Nonlinear Memoryless Devices, With Applications to Quantizers, TR-67-2-BF, Data Systems Division, Litton Systems, Inc., Waltham, MA, 1967.

INITIAL DISTRIBUTION LIST

Addressee	No. of Copies
DARPA, (CAPT H. Cox)	1
NAV SURFACE WEAPONS CENTER, WHITE OAK LABORATORY	1
DWTNSRDC CARD	1
NRL	1
NRL, USRD	1
NORDA	1
NAVELECSYSCOM, PME-124	1
NAVSEASYSCOM, SEA-06H1, -06H1-1, -06H2 (T. Oliver), -09G32(4), -660	8
NAVOCEANSYSCEN (D. Davison)	1
NAVPGSCOL	1
DDC, ALEXANDRIA	12
ENGINEERING SOCIETIES LIBRARY, United Engineering Center, NY 10017	1
C. Veitch, MAR, Inc., Niantic, CT	1
SACLANT ASW Research Centre	1